

Bayesian Post-Selection Inference in the Linear Model

Snigdha Panigrahi* Jonathan Taylor Asaf Weinstein

Stanford University

Abstract

We provide Bayesian inference for a linear model selected after observing the data. Adopting [Yekutieli \(2012\)](#)’s ideas, the Bayesian model consists of a prior and a truncated likelihood. The resulting posterior distribution, unlike in the setup usually considered when performing Bayesian variable selection, is affected by the very fact that selection was applied. After proposing an extension of [Yekutieli](#)’s framework to the case of variable selection, we turn to face the computational challenges associated with the adjusted posterior distribution. A major objection is the intractability of the truncated likelihood. At the core of our methods is a convex approximation to the truncated likelihood, which facilitates sampling from the (approximate) adjusted posterior distribution. We demonstrate in simulations that employing the proposed approximation results in Bayesian procedures that are qualitatively similar to those using the exact truncated likelihood.

Our methods are discussed in the context of recent literature on exact post-selection inference after model selection. These recent works focus on hypothesis testing, and capitalize on reductions achieved by conditioning out nuisance parameters. However, the techniques developed in that venue are generally less appropriate for addressing other questions, like point estimation. On the other hand, relying on an approximation to the full truncated likelihood, the tools we develop allow for more versatility. For example, replacing the genuine truncated likelihood by its approximation, we can approximate the maximum-likelihood estimate as the MAP estimate corresponding to a constant prior. We provide more examples in which our approximation can be employed to address frequentist questions that have not been resolved in existing work on exact post-selection inference.

1 Introduction

An increasing concern about reproducibility in scientific research has recently brought much attention to the problem of selective inference. Informally, the term “selective inference” (sometimes also called *data snooping* or *adaptive data analysis*) refers to a situation in which the analyst interacts with the data to decide what questions about an underlying population she would like to address. In this situation, methods which do not take selection into account are generally invalid: the relevant sampling distribution for inference is altered by selection

*To whom correspondence should be addressed (E- mail: snigdha@stanford.edu).

if the same data that was used to choose the question(s) of interest, is used for inference. For illustration, suppose that an analyst would like to model the relationship between a certain response and a set of potential explanatory variables. Looking at the data, she includes in the model only the few variables which appear to be most correlated with the response. If inference about the chosen model ignores selection, it would probably suggest (by construction) that the underlying selected correlations are large, even if all of them are in fact very small.

The problem of correcting for selection effect is not new, and has been formulated in various ways and with different goals in many previous works. Correspondingly, different approaches to adjusting inference for selection have been suggested, typically with the intention of constructing statistical procedures that have the same (or approximately the same) desirable properties as in the non-selective case. This includes bias-reducing methods based on an extended definition of Uniform Minimum-Variance Unbiased estimation (Robbins, 1988; Cohen and Sackrowitz, 1989), Bayesian approaches (Efron, 2011) and bootstrapping (Simon and Simon, 2013); as well as methods that work by performing simultaneous inference (Berk et al., 2013). Recently, tools from information theory (Russo and Zou, 2015) and Differential Privacy (Dwork et al., 2014, 2015) have also been proposed to quantify and control the effect of selection.

Another common approach is to *condition* on selection, namely, base inference on the likelihood of the observed data when truncated to the set of all realizations that would result in the analyst posing the same question. A conditional approach has been pursued by several authors that addressed selective inference in so-called large-scale inference problems (Efron, 2012). Underlying such problems is a sequence model, where each observation corresponds to a single parameter, and the parameters are typically not assumed to have any relationship with one another. Conditional inference for the effects corresponding to the K largest statistics in the sequence model (as described above for $K = 1$) was proposed in Reid et al. (2014). Zöllner and Pritchard (2007) and Zhong and Prentice (2008) suggested point estimators and confidence intervals for the parameter of a univariate Gaussian distribution conditional on exceeding a fixed threshold, in the context of genome-wide association studies; Weinstein et al. (2013) constructed confidence intervals, and Benjamini and Meir (2014) explored point estimators, for the univariate truncated Gaussian problem with a more flexible (random) choice of cutoff; Yekutieli (2012) proposed to base inference on the truncated likelihood while incorporating a prior on the parameters; and Simonsohn et al. (2014) proposed a frequentist method to assess effect sizes from the distribution of the p-values corresponding to only the selected.

More recently, the practicability of the conditional approach has been extended considerably from the sequence model to the realm of linear models and GLMs (Lee et al., 2013; Taylor et al., 2013, 2014; Lee and Taylor, 2014; Fithian et al., 2014; Tian and Taylor, 2015, among others). In these works the selection protocol is assumed to partition the sample space into polyhedral (or at least convex) sets, fitting many popular ‘automatic’ model selection procedures such as marginal screening, Lasso, forward-selection etc. Other forms of selection rules was considered in Loftus and Taylor (2015); Loftus (2015); Yang et al. (2016).

The techniques developed in this line of work have practical importance as they allow to carry out exact inference after variable selection, which is one of the most popular situations where the problem of selective inference arises.

In the current work, we propose to give post-selection inference in the linear model by using the truncated likelihood in a Bayesian framework. As conceptualized by [George and Yekutieli \(2012\)](#), Bayesian inference after selection relies on a model that prepends a prior to the truncated likelihood. This results in a posterior distribution that is affected by selection, unlike in the usual framework for Bayesian variable selection. While the ideas have already been proposed, [George and Yekutieli](#) have considered only simple selection rules, where computational complications due to the fact that the truncated likelihood replaces the usual likelihood, can be easily overcome. In the current work we tackle the computational difficulties associated with carrying out selection-adjusted Bayesian inference in a more general setup, namely, when the selection rule partitions the sample space into polyhedral subsets.

The main contribution of the current paper is in offering technical tools that enable to provide selection-adjusted Bayesian inference in the linear model, and, importantly, are scalable to work under fairly general selection schemes. The crucial element in our methodology is a tractable approximation to the selection probability given in (10), that has several favorable properties:

- The approximation to (10) leads to a convex approximation to the truncated likelihood, which makes it computationally easy to numerically optimize and to analyze
- The approximation to (10) that we offer is shown in Section 7 to be a large deviations approximation to the exact selection probability
- We demonstrate that when the corresponding approximate maximum likelihood estimator (MLE) is not a consistent estimator, then applying similar methods to a randomized version of the data can produce a consistent estimator.

That the truncated likelihood is involved connects our work to existing frequentist work on post model-selection inference. However, performing Bayesian inference requires new techniques: the frequentist theory developed so far for the linear model (e.g., [Lee et al., 2013](#)) focuses on hypothesis testing for linear functionals of the parameter, and relies on reducing the problem to a truncated problem free of nuisance parameters. By contrast, in our framework inference relies on the selection-adjusted posterior, which involves the *full* truncated likelihood. Consequently, even inference for a one-dimensional linear projection of the parameter vector requires to handle nuisance parameters. Moreover, the frequentist tests and confidence intervals developed in (e.g., [Lee et al., 2013](#)), based on a univariate pivot statistic, have some optimality properties under a particular generative model for the data (referred to as the “saturated model” below; see [Fithian et al., 2014](#)) but not in general (for example, under the selected model they do not). Our procedures, on the other hand, are not strongly tied to a particular generative model: we can construct them for essentially any generative model by sampling from an approximate posterior distribution (see Section 5).

Bayesian inference based on the truncated likelihood has been suggested before by [Bayarri and DeGroot \(1987\)](#)¹ who used it in situations where the statistician has access to truncated data in the first place; for example, inference on the distribution of heights in a certain period in the past might have to be based on historical records of the heights of members in the army, for which there was a minimum required height ([Wachter and Trussell, 1982](#)). Modeling selection-adjusted inference as a truncation problem, [Yekutieli \(2012\)](#) used the truncated likelihood along with a prior to give Bayesian inference on selected parameters in the sequence model. The situation considered in the current paper, similarly to the focus in [Yekutieli \(2012\)](#), is where truncation is imposed by the statistician herself when adaptively choosing the target of inference.

When providing adjusted Bayesian inference, we emphasize the advantages of using a *randomized* response for selection, as proposed by [Tian and Taylor \(2015\)](#). Randomizing before applying selection may seem odd because it can (and often does) result in selecting a different set of variables, but so does the commonly used practice of sample splitting (in fact, if the split is random, then, as pointed out in [Tian and Taylor \(2015\)](#), sample splitting is an example of using a randomized response). On the up side, randomization is a way to ensure that more information is reserved for the post-selection stage. This usually results in improved accuracy of post-selection inference procedures, for example shorter credible intervals.

The rest of the paper is organized as follows. In Section 2 we describe a framework for Bayesian selection-adjusted inference in the linear model. We motivate the use of the proposed framework and contrast it with the standard Bayesian variable-selection framework in Section 3. Section 4, the centerpiece of this paper, proposes an approximation to the selection-adjusted likelihood, which we use in Section 5 to provide adjusted approximate inference for selected parameters in a simulation example. Section 6 is devoted to point estimation, where we study the approximate maximum a-posteriori (MAP) estimator. We provide some theoretical support for using the proposed approximation in Section 7 and prove consistency guarantees of the selective MLE, obtained by maximizing the pseudo selection-adjusted likelihood, when selection is based on a randomized response. The supplement includes an analysis of real data, where we also compare our methods to existing frequentist methods for post-selection inference in the linear model and demonstration of a few useful applications of the pseudo truncated likelihood in frequentist inference.

2 Bayesian selection-adjusted inference in the linear model

This section presents a framework for Bayesian adjusted inference after variable selection in the linear model. [Yekutieli \(2012\)](#) described a general framework for Bayesian adjusted inference, however his framework is aimed more at large-scale inference problems; the case of model selection calls for additional definitions. The main conceptual difference is that, in

¹we thank Ed George for bringing this reference to our attention.

the model selection problem, any parameter selected depends (by definition of a parameter) on the generative model, and, as proposed in [Fithian et al. \(2014\)](#), we want to allow the analyst to choose also the generative model after observing the data. The generative model and the truncating event together determine the truncated likelihood.

We assume that we observe a matrix $X = [X_1 \cdots X_p] \in \mathbb{R}^{n \times p}$ and a random p -dimensional vector

$$Y|X \sim F.$$

X is assumed to be a fixed design matrix, hence we will often write the above just as $Y \sim F$. Informally, selective inference in the linear model consists of two stages:

1. **Selection.** Choose a subset $\widehat{E}(Y) = E \subseteq \{1, \dots, p\}$ as a function of the data (X, Y)
2. **Inference.** Provide inference for

$$\beta^E = \operatorname{argmin}_{b^E \in \mathbb{R}^{|E|}} \|\mathbb{E}(Y) - X_E b^E\|_2^2 = (X_E^T X_E)^{-1} X_E^T \mathbb{E}(Y) \quad (1)$$

where $X_E = [X_j : j \in E] \in \mathbb{R}^{n \times |E|}$ is the matrix consisting of only the columns of X with indices in E , and is assumed to have full rank.

The target of inference, β^E , is sometimes referred to as the *population* least-squares vector corresponding to the sub-model E , and, in words, is the coefficient vector of the best approximation to the mean of Y in the linear space spanned by the columns of X_E . To be formally defined as a parameter, a generative model needs to be specified for Y , that is, a family \mathcal{F} of distributions to which the unknown distribution F is assumed to belong.

In contrast to the traditional (non-selective) statistical framework, we allow the statistician to choose a generative model *after* she looks at the data, in the same way she chooses the *target* of inference only after looking at the data. The minimum assumptions we make for the generative model are that

$$Y \sim N_n(\mu, \Sigma) \quad (2)$$

with $\mu \in \mathbb{R}^n$ unknown and Σ known. If the statistician chooses to proceed with (2) as the generative model, we will say that inference is given under the *saturated model*, and in this case

$$\beta^E = (X^T X)^{-1} X^T \mu.$$

More generally, we allow the statistician to restrict μ in (2) to a linear subspace of \mathbb{R}^n , in which case the generative model is the linear model

$$Y \sim N_n(X^* \beta^*, \Sigma) \quad (3)$$

where X^* may include columns of X . For example, one may take $X^* = X_E$, in which case

$$\beta^E = (X_E^T X_E)^{-1} X_E^T X_E \beta^* = \beta^*$$

and we will say that inference is given under the *selected model*. But the generative model does not have to agree with the selected model (or the saturated model): one may take

$X^* = X_{E'}$ for $E' \neq E$, including the choice $E' = \{1, \dots, p\}$, in which case the saturated model obtains; or even include as columns of X^* predictors that were not originally recorded in X but now – after selecting E – seem relevant. In general, then, inference is given for

$$\beta^E = (X_E^T X_E)^{-1} X_E^T X^* \beta^*.$$

It is important that the specification of X^* depends on Y only through E . So, for example, it is not allowed to run Lasso again, now for X augmented with some other variables that were not originally recorded in X , and let X^* depend on the output. Indeed, this may lead to bias that is not accounted for. We remark that allowing to specify the generative model (X^*) after obtaining E may reduce the difficulty of specifying a prior for the problem, because only a prior for β^* is required rather than a prior on all possible sub-models. Indeed, in some cases domain knowledge could be used to elicit a prior for the parameters of the generative model.

The question of what generative model to use is, unavoidably, somewhat arbitrary, and the analyst has to make some choice. It is worth emphasizing, however, that the choice of the generative model can have important consequences not only on the interpretation of the selected parameter, but also on the properties of statistical procedures and on their implementation, as explained in [Fithian et al. \(2014\)](#). For example, besides power considerations, there is an essential difference between the saturated model, $X^* = I_n$, and the other cases, $\text{rank}(X^*) < n$, in terms of the computational difficulty in constructing optimal (UMPU) tests; see Section 9.2.2 of the Supplementary Material.

Throughout the paper we will restrict attention to selection rules² which partition the sample space into convex polyhedral sets,

$$\{y : A_E y \leq b_E\} \tag{4}$$

for corresponding A_E and b_E associated with E . In previous works ([Lee et al., 2013](#); [Taylor et al., 2014](#), among others), it was shown that for various popular ‘automated’ selection algorithms, the set $\{y : \widehat{E}(y) = E\}$ can indeed be written in the form (4).

Inference for β^E is based on the *selection-adjusted* likelihood

$$f_S(y|\beta^*) = \frac{f(y|\beta^*)}{\mathbb{P}(A_E Y \leq b_E|\beta^*)} I(A_E y \leq b_E), \tag{5}$$

where

$$f(y|\beta^*) = (2\pi)^{-n/2} |\Sigma|^{-1/2} \exp\left(-\frac{1}{2}(y - X^* \beta^*)^T \Sigma^{-1} (y - X^* \beta^*)\right).$$

To obtain a Bayesian model we incorporate a prior

$$\beta^* \sim \pi(\beta^*)$$

²For some popular selection rules like the Lasso, forward stepwise etc., we condition additionally on the signs of the active coefficients to obtain polyhedral regions

into the model given by (5). Formally, then, (β^*, Y) have a *selection-adjusted* joint distribution given by

$$f_S(\beta^*, y) = \pi(\beta^*) f_S(y|\beta^*) = \pi(\beta^*) \frac{f(y|\beta^*)}{\mathbb{P}(A_E Y \leq b_E|\beta^*)} I(A_E y \leq b_E). \quad (6)$$

The framework for adjusted Bayesian inference after model selection is no different than an ordinary Bayesian framework, where the joint distribution of the parameter and the data is given by (6). In particular, if $L(a, \beta^*)$ is a loss function and $\delta_S(y)$ is a rule that maps points in the restricted sample space $\{y : A_E y \leq b_E\}$ to actions, then the adjusted Bayes risk of δ w.r.t. π is

$$r_S(\pi, \delta) = \iint L(\delta_S(y), \beta^*) f_S(\beta^*, y) dy d\beta^*.$$

Correspondingly, the Bayes rule minimizes the selection-adjusted posterior expected loss,

$$\rho_S(\pi, \delta_S, y) = \int L(\delta_S(y), \beta^*) \pi_S(\beta^*|y) d\beta^*$$

where

$$\pi_S(\beta^*|y) = \frac{f_S(\beta^*, y)}{\int f_S(\beta^*, y) d\beta^*} \propto \pi(\beta^*) f_S(y|\beta^*) \quad (7)$$

is the selection-adjusted posterior distribution. Note that this can be easily adapted for inference on a function $t(\beta^*)$ by taking $L(a, \beta^*) = L_t(a, t(\beta^*))$ where L_t is a loss function defined on an action set and state set chosen for $t(\beta^*)$. This is just a formal way of saying that the posterior distribution of β^* determines the optimal procedures for any function of β^* ; specifically, this is true for β^E , the function of β^* which inference is required for. In later sections we base post-model selection inference on the (approximate) selection-adjusted posterior distribution (7) of β^* , without necessarily committing to a formal decision theoretic framework.

3 Motivation and comparison with standard Bayesian model uncertainty framework

Our aim in this section is to advocate the use of a truncated likelihood in place of the usual likelihood when providing Bayesian inference after model selection. Following Yekutieli (2012), we model Bayesian selection-adjusted inference as a truncation problem. Namely, the fact that inference for a specific parameter is given only if an associated event in the observation space occurs, elicits a truncated likelihood for this parameter. As shown by Yekutieli (2012), the fact that a truncated likelihood is involved, does not immediately imply that selection should alter posterior inference. In fact, the common perception is that a Bayesian is indifferent to selection: the usual formalism suggests that $\pi(\theta|y, Y \in S) = \pi(\theta|y)$ for $y \in S$, hence no adjustment for selection seems necessary when conditioning on y . Yekutieli (2012) demonstrated that this argument is valid when truncation is applied to the

pair $(\theta, Y) \sim \pi(\theta)f(y|\theta)$, but not when truncation is applied to Y conditionally on $\theta \sim \pi$. To illustrate the difference between the two situations, consider the following example, inspired by [Mandel and Rinott \(2007\)](#).

Example 1. *There is an infinite number of bags of coins with different probabilities for a Head. The composition of each of the bags is the same, and such that if a coin is randomly selected, then the probability of landing on a Head has distribution π . We are interested in constructing a posterior interval for P , the probability of Head, in the following two experiments.*

Situation I: Draw a coin at random from the first bag, and toss it 10 times. If at least 3 Heads observed, report $Y = \#$ of Heads observed. If not, move on to the second bag, draw a coin at random, and toss it n times. If at least 3 Heads observed, report $Y = \#$ of Heads observed. If not, move on to the third bag, and so on until the first time that 3 Heads or more are observed. Report $Y = \#$ of Heads observed, and let P be the (unobserved) probability of Head for the last coin drawn.

Situation II: Draw a coin at random from the first bag. Let P be the (unobserved) probability of Head for the coin drawn. Toss the coin 10 times. If at least 3 Heads observed, report $Y = \#$ of Heads observed. If not, toss the same coin again 10 times. If at least 3 Heads observed, report $Y = \#$ of Heads observed. If not, toss the same coin again 10 times. Keep going until the first time that 3 Heads or more are observed, and report $Y = \#$ of Heads observed.

[Yekutieli \(2012\)](#) referred to Situation I as the “random” parameter case and to Situation II as the “fixed” parameter case. In both cases, the number of Heads is reported only if at least 3 Heads are observed; therefore, the likelihood in both cases is a truncated likelihood, say $f_S(y|p)$, where $f_S(y|p) = f(y|p)I_{\{y \geq 3\}}(y)/\mathbb{P}(Y \geq 3|p)$. Yet selection matters only in Situation II. Indeed, in this case the selection-adjusted posterior is $\pi_S(p|y) \propto \pi(p)f_S(y|p)$, whereas in Situation I, the selection-adjusted posterior is

$$\pi_S(p|y) \propto \pi(p|Y \geq 5)f_S(y|p) = \frac{\pi(p)f(Y \geq 5|p)}{\mathbb{P}(Y \geq 5)} \times \frac{f(y|p)I_{\{y \geq 5\}}(y)}{\mathbb{P}(Y \geq 5|p)} \propto \pi(p)f(y|p).$$

Note that in Situation I two sources contribute to “shrinking” the likelihood: one is the prior $\pi(p)$, and the other is the term $1/\mathbb{P}(Y \geq 5|p)$, that gives a-priori preference to values of p under which the selection event $\{Y \geq 5\}$ is less likely to occur.

As recommended by [Yekutieli \(2012\)](#), choosing to model θ as a “random” parameter or a “fixed” parameter should depend on the context and be done case by case. [Bayarri and DeGroot \(1987\)](#) proposed to analyze results reported in a scientific journal that has a policy of publishing only statistically significant findings, when essentially treating the underlying effect size of a published result as a “fixed” parameter. Indeed, their perspective reflects a model in which the pair (θ_i, Y_i) for a published result is generated by sampling $\theta_i \sim \pi$ once, and then sampling $Y_i|\theta_i \sim f(y_i|\theta_i)$ repeatedly, keeping the first Y_i which passes a significance threshold. We argue that the “random” parameter model is more suitable for this example. That is, to model a randomly selected reported result in the journal as being generated by repeatedly sampling a pair $(\theta_i, Y_i) \sim \pi(\theta_i)f(y_i|\theta_i)$ and keeping the first pair for

which Y_i passes the significance threshold. This model is justifiable because it is natural to think of selection as acting on θ_i (as well as on Y_i , of course) in this case. In other words, reported results tend to come with larger effect sizes. Note that the resulting analyses will be substantially different, because the former entails appending the truncated likelihood $f_S(y|\theta)$ to $\pi(\theta)$, whereas the latter entails appending the unadjusted likelihood $f(y|\theta)$ to $\pi(\theta)$.

On the other hand, for inference after model (or variable) selection, we argue that treating the parameter β^* as a “fixed” parameter makes more sense. From our perspective, the parameter β^* and the data Y are not *actually* jointly distributed; instead, only the data is really random, while the prior on β is used to reflect prior beliefs on an “unknown constant”, to use the words of Yekutieli (2012).

To try and convince the reader that treating β^* as “fixed” rather than “random” parameter might be a good idea, we follow Example 2 of Yekutieli (2012) and compare the coverage of credible intervals constructed under a non-informative prior in the linear model, when treating the parameter as “fixed” and when treating the parameter as “random”.

Example 2. We fix an $n \times p$ design matrix X with $n = 100, p = 50$, and in each of 100 rounds draw a pair (β, Y) , where the components of $\beta \in \mathbb{R}^{50 \times 1}$ are drawn i.i.d. from

$$\pi(\beta) = 0.9 \cdot \mathcal{N}(\beta; 0, 0.1) + 0.1 \cdot \mathcal{N}(\beta; 0, 3), \quad (8)$$

a mixture of two zero-mean normal distributions, one with small variance 0.1 and the other with larger variance 3, and $Y|\beta \sim N_n(X\beta, I)$. We run the Lasso with a fixed $\lambda = 1.56$ and denote by $E = \hat{E}(Y) \subseteq \{1, \dots, 50\}$ the set corresponding to nonzero Lasso estimates. We then construct adjusted and unadjusted marginal posterior intervals for the components of $\beta^E = (X_E^T X_E)^{-1} X_E^T X \beta$ assuming a noninformative prior $\pi(\beta) \propto 1$ instead of the true prior, which is regarded as unknown. The first interval is constructed using the unadjusted Normal likelihood for β . The second interval is constructed using the (approximate) selection-adjusted likelihood and taking $X^* = X_E$,

$$f_S(y|\beta) \propto \frac{\exp\left(-\frac{1}{2}\|y - X_E \beta\|_2^2\right)}{\tilde{\mathbb{P}}(\hat{E}(Y) = E|\beta)} I_{\{y: \hat{E}(y) = E\}}(y). \quad (9)$$

We also construct randomized selection-adjusted intervals and data-carved selection-adjusted intervals, see Section 5 for details. Above, $\tilde{\mathbb{P}}(\hat{E}(Y) = E|\beta)$ is our barrier approximation to the exact selection-adjusted likelihood, discussed in Section 4. Because the posterior distribution of β^E under the noninformative prior is $N_{|E|}(\hat{\beta}^E, (X_E^T X_E)^{-1})$, where $\hat{\beta}^E = (X_E^T X_E)^{-1} X_E^T Y$ is the LS estimate for β^E , the equal-tails 0.95 unadjusted marginal posterior interval for β_j^E is $\hat{\beta}_j^E \pm \sigma_j(E) \cdot z_{1-0.05/2}$ where $\sigma_j^2(E) = [(X_E^T X_E)^{-1}]_{jj}$. The selection-adjusted marginal credible intervals are computed using posterior sampling, as described in Section 5.

Without selection the unadjusted intervals for $\beta_j, j = 1, \dots, 50$ are expected to have the property that the (Bayesian) FCR, the expectation of the proportion of noncovering constructed intervals, is controlled at level 0.95, because these are valid frequentist marginal

intervals. Indeed, with no selection (i.e., when constructing intervals for all 50 components of β), the unadjusted intervals had observed FCR of 0.951. But with selection applied as described above, the observed FCR was only 0.75. On the other hand, the selection-adjusted posterior intervals had observed FCR of 0.85 – a considerable improvement over the unadjusted intervals – despite using a misspecified likelihood (because we computed them using $X^* = X_E$ in the generative model, while the true generative model has $X^* = X$) and an approximation instead of the exact term in the denominator of (9). With randomization (see Section 5.2) the observed FCR increases to 0.94, and with data-carving (see Section 5.1) the observed FCR is 0.942. The improved robustness of the various selection-adjusted intervals implies that these methods tend to have better frequentist coverage than the unadjusted intervals. This example also demonstrates the potential benefits of randomization and data-carving, ensuring better frequentist properties of the selection-adjusted intervals.

4 Approximations to the truncated likelihood

From now on, let $\|\cdot\|$ denote the ℓ_2 norm, unless specified otherwise. To be able to take advantage of Bayesian computational methods, an immediate challenge that presents itself is to evaluate the selection-adjusted likelihood (5) as a function of β^* . The difficulty is in computing the *adjustment factor*

$$\mathbb{P}(A_E Y \leq b_E | \beta^*) \quad (10)$$

as a function of β^* for given matrix A_E and vector b_E . Since (10) in general has no closed form, we will propose an approximation of the adjustment factor that is computationally tractable while exhibiting similar properties as the exact expression. We will assume from now on, unless otherwise noted, that $\Sigma = \sigma^2 I_n$. Extending our methods to the case of general (known) covariance is simple, especially since for (10) we can assume $\text{cov}(Y) = I_n$ without loss of generality; indeed,

$$\mathbb{P}(AY \leq b | \beta^*) = \mathbb{P}(\tilde{A}\tilde{Y} \leq b | \beta^*)$$

with $\tilde{A} = A\Sigma^{1/2}$, $\tilde{Y} = \Sigma^{-1/2}Y \sim N_n(\Sigma^{-1/2}X^*\beta^*, I_n)$.

Before delving into the details, we would like to convince the reader that evaluating the adjustment factor is essential. The approach relied upon in recent frequentist work on exact post-selection inference, was able to bypass handling the adjustment factor directly. We explain why this technique is not suitable in our case. Consider inference under the saturated model (i.e., $X^* = I_n$; $\beta^* \in \mathbb{R}^n$), $Y \sim N_n(\beta^*, I_n)$, and suppose that we want to test a null hypothesis regarding a one-dimensional projection of β^* , say $\nu = \eta^T \beta^*$, conditional on $Y \in \{y : Ay \leq b\}$. Decomposing

$$Y = (\eta^T Y) \frac{\eta}{\|\eta\|^2} + Z$$

where $Z = (I - P_\eta)Y$ is the projection of Y orthogonally to η , a crucial observation is that the conditional distribution of $\eta^T Y$ given that $Y \in \{y : Ay \leq b\}$ and $Z = z$, is just

the distribution of $\eta^T Y$ truncated to the interval $[\gamma^-(z), \gamma^+(z)]$ for known functions γ^+, γ^- associated with A, b . Importantly, this distribution depends only on the scalar parameter of interest, $\nu = \eta^T \beta^*$. This result is key to construction of frequentist tests (and, by duality, of confidence intervals) which are valid conditionally, and hence also unconditionally, on Z . In fact, [Fithian et al. \(2014\)](#) show that under the saturated model, Uniformly Most Powerful Unbiased (UMPU) tests for testing a two-sided hypothesis are of this form; if instead of a saturated, the model used for inference has X^* with $\text{rank}(X^*) < n$, constructing optimal test again requires to handle terms like (10); this is generally the case when inference under the *selected* model is of interest (in Section 9.2.2 of the supplement we propose to employ our methods for this task).

The situation is therefore congenial when performing selective hypothesis testing. But it is more complicated when general likelihood-based inference is considered, estimation being a good example. To see that the conditioning to eliminate nuisance parameters does not lead to a reduction in the Bayesian setting, we may write the selection-adjusted likelihood as

$$f_S(\eta^T y, z | \beta^*) = f_S(\eta^T y | z, \beta^*) f_S(z | \beta^*) = g_{\nu, \sigma^2}^{[\gamma^-(z), \gamma^+(z)]}(\eta^T y) f_S(z | \beta^*).$$

where $g_{\nu, \sigma^2}^{[\gamma^-(z), \gamma^+(z)]}$ is the density of a $N(\nu, \sigma^2)$ variable truncated to $[\gamma^-(z), \gamma^+(z)]$. For fixed z , if $f_S(z | \beta^*)$, the conditional density of a lower dimensional projection of Y , did not depend on ν , then computing the posterior density of ν would conveniently reduce to computing the posterior under a univariate truncated normal likelihood, a case that was studied extensively by [Yekutieli \(2012\)](#). Unfortunately, $f_S(z | \beta^*)$ generally *does* depend on ν , and so conditioning on z is not particularly helpful.

Now that we have motivated the need for it, we turn to propose a workable approximation for the adjustment factor. Note that at this point the discussion revolves around the selection-adjusted likelihood only, therefore any of the approximations that we will suggest for the truncated likelihood can be used for frequentist instead of Bayesian inference.

4.1 A first attempt

The quantity of interest is the probability (10) under $Y \sim N_n(X^* \beta^*, \sigma^2 I)$. We already noted that the exact expression is in general a complicated function of β^* ; still, it is clear that this function should take on larger values when the mean $X^* \beta^*$ lies closer to the polytope $\mathcal{K} = \{y : A_E y \leq b_E\}$. It would therefore be sensible to consider the approximation

$$\mathbb{P}(A_E y \leq b_E | \beta^*) \approx \exp \left(- \frac{\|X^* \beta^* - P_{\mathcal{K}}(X^* \beta^*)\|^2}{2\sigma^2} \right) \quad (11)$$

where $P_{\mathcal{K}}(X^* \beta^*)$ is the projection of $X^* \beta^*$ onto \mathcal{K} . Indeed, the following theorem shows that the right-hand side of (11) can be derived as an upper bound for the exact term by generalizing a Chernoff bound. In fact, the right hand side of (11) is a large deviations approximation to the left hand side. We discuss this in detail in Section 7.

Theorem 4.1 (Chernoff approximation). *Let $Y \sim N_n(X^*\beta^*, \sigma^2 I_n)$ and let $\mathcal{K} \subseteq \mathbb{R}^n$ be convex and compact. Then*

$$\mathbb{P}(Y \in \mathcal{K} | \beta^*) \leq \exp \left(-\frac{\|X^*\beta^* - \mathcal{P}_{\mathcal{K}}(X^*\beta^*)\|^2}{2\sigma^2} \right)$$

where

$$\mathcal{P}_{\mathcal{K}}(\mu) = \inf \{ \|\mu - z\|^2 : z \in \mathcal{K} \}$$

is the projection of mean vector $\mu = X^*\beta^*$ onto \mathcal{K} .

Remark 1 (Relaxing compactness). *The polyhedral set corresponding to many popular variable selection rules (as in our examples) is not compact. While compactness is required for our proof, one should be able to remove this condition by considering, for a fixed π , a sufficiently large compact, convex subset of \mathcal{K} that would work for all β^* in a bounded set of probability close to 1 under π . In any case, the purpose of Theorem 4.1 is only to motivate the proposed approximation.*

Remark 2. *While the theorem implies that the Chernoff approximation is sensible for general convex sets, the computation of the approximation is less burdensome when the set \mathcal{K} is a polyhedral set: finding $\mathcal{P}_{\mathcal{K}}(\mu)$ is then a convex optimization problem, and there are available tools to solve it. Specifically, in the appendix we describe how to employ an algorithm based on ADMM (Boyd et al., 2011) to numerically solve the optimization problem.*

4.2 Approximation with a barrier function

While the Chernoff approximation (11) exhibits desirable properties, it may be too crude for certain purposes. Indeed, it incurs no penalty when $X^*\beta^*$ is inside the selection region, even though the exact probability (10) is not constant on $\{\beta^* : X^*\beta^* \in \mathcal{K}\}$. In Section 6 we discuss point estimation and show that using the approximation (11) produces no shrinkage, while we *do* expect the point estimates to have the property of shrinking near the selection boundary. This motivates looking for a refined approximation for the adjustment factor.

In general, consider an approximation of the form

$$\mathbb{P}(A_E y \leq b_E | \beta^*) \approx \exp(-h_\sigma(X^*\beta^*)). \quad (12)$$

Then the approximation in (11) has

$$h_\sigma(X^*\beta^*) = \frac{1}{2\sigma^2} \|X^*\beta^* - \mathcal{P}_{\mathcal{K}}(X^*\beta^*)\|^2 = \inf_{z \in \mathcal{K}} \frac{1}{2\sigma^2} \|X^*\beta^* - z\|^2. \quad (13)$$

Note that this can be rewritten as

$$\inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2\sigma^2} \|X^*\beta^* - z\|^2 + I_{\mathcal{K}}(z) \right\} \quad (14)$$

where

$$I_{\mathcal{K}}(z) = \begin{cases} 0 & z \in \mathcal{K} \\ \infty & \text{otherwise} \end{cases} \quad (15)$$

is an indicator function for the set \mathcal{K} .

More generally, consider

$$h_{\sigma}(X^*\beta^*) = \inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2\sigma^2} \|X^*\beta^* - z\|^2 + w_{\sigma}(z) \right\} \quad (16)$$

where $w_{\sigma}(z)$ is a non-negative function that should, informally, reflect a preference for values of z farther away from the boundary and inside the selection region; in other words, $w_{\sigma}(z)$ should take on smaller values for such z . Throughout, when $\sigma^2 = 1$, we write simply $w(z)$ instead of $w_1(z)$ and $h(z)$ instead of $h_1(z)$.

Refining the choice in (15), we propose to substitute for $w_{\sigma}(z)$ a barrier function,

$$w_{\sigma}^{\text{bar}}(z) = \begin{cases} \sum_{i=1}^m \log \left(1 + \frac{\sigma}{b_i - a_i^T z} \right) & z \in \mathcal{K} \\ \infty & \text{otherwise} \end{cases}. \quad (17)$$

Correspondingly, (10) will be approximated as

$$\exp(-h_{\sigma}^{\text{bar}}(X^*\beta^*))$$

with

$$h_{\sigma}^{\text{bar}}(\beta^*) = \inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2\sigma^2} \|X^*\beta^* - z\|^2 + w_{\sigma}^{\text{bar}}(z) \right\}. \quad (18)$$

Note that as $z \in \mathcal{K}$ approaches the boundary of the selection region, $\sum_{i=1}^m \log(1 + \sigma/(b_i - a_i^T z))$ will tend to ∞ , which is typical of a barrier function. Hence, for fixed $X^*\beta^*$, (18) encourages z to lie deeper inside the selection region, where $b_i - a_i^T z$ take on larger values. Another advantage of using (18) instead of (14) is that the optimization problem is now not constrained.

The left panel of Figure 1 displays the two approximations to the adjustment factor as well as the exact expression (using Mills' approximation for large $|\mu|$) for the truncated univariate example. It can be seen that the Barrier approximation is more accurate than the Chernoff approximation. As we will see in Section 6, the modification above will also have a desirable shrinking effect on our point estimates. Corollary 7.2 in Section 7 says that the barrier modification also yields a large deviation approximation to the adjusted likelihood as long as the barrier function (with scale $1/n$) converges pointwise to the discrete indicator penalty $I_{\mathcal{K}}$.

5 Inference based on the selection-adjusted posterior

Equipped with a tractable approximation to the adjustment factor, we can now employ generic Bayesian computational methods to give selection-adjusted Bayesian inference for




Figure 1: LEFT: Approximations to $\log \mathbb{P}(A_E Y \leq b_E | \beta^*)$ in the univariate case, $Y \sim N(\beta^*, 1)$. Here $\{A_E Y \leq b_E\} = \{Y > 0\}$. The Barrier approximation is closer to the exact function.

RIGHT: Selective MLE with Barrier approximation for the univariate example, $Y \sim N(\beta^*, 1)$ truncated to $\{y > 0\}$. The solid blue line is the approximate selective MLE under randomized response, $Y^* = Y + W$ truncated to $\{Y^* > 0\}$ (note that this estimator is defined for $y \in \mathbb{R}$). All other estimates are for the standard (non-randomized) setting. The straight line is the identity, corresponding to the unadjusted estimate.

the parameters of the model. In general, we will use a diffuse prior for β^* and provide selection-adjusted inference for scalar functions of the parameter, $t(\beta^*)$, $t: \mathbb{R}^k \rightarrow \mathbb{R}$. An example is $\beta_j^E = e_j^T (X_E^T X_E)^{-1} X_E^T X^* \beta^*$. When providing inference based on draws from the (approximate) selection-adjusted posterior, we use the selected model, $X^* = X_E$, as the generative model. In principle, if the general model $y \sim N(X^* \beta^*, \sigma^2 I)$ is preferred, one would draw posterior samples of β^* and then compute the projection $\beta^E = (X_E^T X_E)^{-1} X_E^T X^* \beta^*$, for each sampled value of β^* .

We use a Langevin random walk to obtain a sample of size n_S from the (approximate) selection-adjusted posterior, with the $(K + 1)$ -th iterate drawn as

$$\beta_{(K+1)}^* = \beta_{(K)}^* + \gamma \nabla \log \pi_S(\beta_{(K)}^* | y) + \sqrt{2\gamma} \epsilon_{(K)}$$

where $\epsilon_{(K)}$ for $K = 1, 2, \dots, n_S$ are independent draws from a centered gaussian with unit variance and γ is a suitably chosen step size. This sampler takes a noisy step along the gradient of the log-posterior, with no accept-reject step unlike the usual Metropolis Hastings (MH) algorithm. Hence, at each draw of the sampler, we compute the gradient of the log selection-adjusted posterior based on the optimizer in the optimization problem associated with (18), see (24). Alternatively, we could employ a MH algorithm to compute $h_\sigma^{\text{bar}}(X^* \beta^*)$;

that would require the value of optimization problem in (18) for computing the acceptance ratio each time. When approximating the adjustment factor, (13) could be used instead of (18); however, the previous section suggests that the barrier approximation produces more accurate results.

5.1 Inference employing Data-carving

Fithian et al. (2014) proposed a general scheme, *data-carving*, to increase precision of post-selection inference. In that scheme a small portion of the data is reserved for the inference stage, i.e., not used for selection; while the entire data is used for inference conditional on selection. Thus, as opposed to data splitting, which ignores the hold-out portion of the data completely at the inference stage, data carving discards only the information in the hold-out portion that is captured by the random variable $1_{A_q}(Y)$. Fithian et al. (2014) show that, as expected, selective tests based on data-carving can be substantially more powerful than selective tests based on data splitting. Data carving can certainly be incorporated into our Bayesian framework. Let the data be

$$Y = \begin{bmatrix} Y^{(1)} \\ Y^{(2)} \end{bmatrix} \sim N \left(\begin{bmatrix} X^{*(1)}\beta^* \\ X^{*(2)}\beta^* \end{bmatrix}, \sigma^2 \begin{bmatrix} I_{n_1} & 0 \\ 0 & I_{n_2} \end{bmatrix} \right).$$

The selection event is $\{A_E^{(1)}Y^{(1)} \leq b_m\}$ where $A_E^{(1)} = A_E S$, $S = [I_{n_1} \ 0_{(n-n_1) \times n_1}]^T$.

Since the selection event is determined by $Y^{(1)}$ only, the selection-adjusted joint distribution of the parameter and the data is

$$f_S(\beta^*, y^{(1)}, y^{(2)}) = \pi(\beta^*) f_S(y^{(1)}|\beta^*) f(y^{(2)}|\beta^*). \quad (19)$$

Note that, rearranging the terms on the right hand side and using Bayes rule, we have

$$f_S(\beta^*, y^{(1)}, y^{(2)}) \propto \pi(\beta^*|y^{(2)}) f_S(y^{(1)}|\beta^*).$$

That is, with data carving, the posterior distribution can be interpreted as arising from updating first the distribution of β^* based on the untruncated likelihood of $Y^{(2)}$ only, and using this as the updated prior for β^* in a selective problem involving only $Y^{(1)}$.

5.2 Inference employing randomization

Data-carving may reduce the length of confidence intervals because it guarantees that the information for the parameter in the hold-out portion, is not affected by selection. Another general strategy for ensuring that enough information is reserved for inference, suggested by Tian and Taylor (2015), is to introduce user-generated randomness in the selection stage. Indeed, randomization tends to smooth out the effect of truncation, which results in more *left-over information* (introduced in Fithian et al. 2014, this is the expectation of the conditional fisher information given the indicator $1_{A_q}(Y)$). In fact, as they pointed out, data-splitting and data-carving that use random splitting, are particular schemes for selective inference

with a randomized response. [Tian and Taylor \(2015\)](#) proposed a specific *Gaussian randomization* scheme, in which the observed data is still the vector y , but truncation is applied to $Y^* = Y + W$ where $W \sim N_n(0, \gamma^2 I)$ and independent of Y . If the selection region remains $\{y^* : A_E y^* \leq b_E\}$, then the selection-adjusted joint distribution of (β^*, y, w) is now

$$f_S(\beta^*, y, w) = \pi(\beta^*) \frac{f(y, w|\beta^*)}{\mathbb{P}(A_E(Y + W) \leq b_E|\beta^*)} I(A_E(y + w) \leq b_E).$$

Integrating out w gives the selection-adjusted joint distribution of β^* and the observed data y as

$$\begin{aligned} f_S(\beta^*, y) &= \pi(\beta^*) \frac{f(y|\beta^*)}{\mathbb{P}(A_E(Y + W) \leq b_E|\beta^*)} \mathbb{P}(A_E(y + W) \leq b_E) \\ &\propto \pi(\beta^*) \frac{f(y|\beta^*)}{\mathbb{P}(A_E(Y + W) \leq b_E|\beta^*)}. \end{aligned} \tag{20}$$

In the last step we used the fact that W is independent of Y and its distribution does not depend on β^* . To provide adjusted Bayesian inference in this situation, we can proceed as in the non-randomized setup after we obtain an approximation for $\mathbb{P}(A_E(Y + W) \leq b_E|\beta^*)$; note that this is equal to $\mathbb{P}(A_E \tilde{Y} \leq b_E|\beta^*)$ where $\tilde{Y} \stackrel{d}{=} Y + W \sim N_n(X^* \beta^*, (\sigma^2 + \gamma^2)I)$, hence reduces to the computation of the adjustment factor when σ^2 is replaced by $\sigma^2 + \gamma^2$.

Example 2 (continued). We return to *Example 2 of Section 3* and obtain credible intervals and point estimates for the components of $\hat{\beta}^E$. Recall that $Y|\beta \sim N_n(X\beta, I)$ but we assume the selected model $Y|\beta^E \sim N_{|E|}(X_E \beta^E, I)$ when providing inference. Also, a noninformative prior is used for β^E instead of the true prior (8). [Figure 3](#) displays selection-adjusted as well as unadjusted credible intervals and point estimates for a particular realization $Y = y$: the left panel is for the non-randomized and the right panel is for the randomized setup. In applying the randomized methods, selection is applied to a perturbed version of y , $Y^* = y + W$ with $W \sim N_n(0, \gamma^2 I)$, $\gamma^2 = 0.1$. This resulted in selecting a different set E (from the figure we can tell that $|E| = 6$ without randomization, and $|E| = 7$ with randomized response).

The adjusted intervals are considerably longer than the unadjusted intervals, and are usually not symmetric about the estimate (black filled circles in the figure). The figure also shows that randomization can contribute significantly to shortening the credible intervals, which illustrates the gain in power with randomization. The adjusted-MAP point estimates, discussed in the next section, seem to shrink the estimates $\hat{\beta}_j^E$ toward zero in the left panel of the figure. the adjusted mean point estimates and the equi-tailed adjusted credible intervals are computed using the barrier approximation and posterior sampling.




Figure 2: Selection-adjusted credible intervals and point estimates for selection via Lasso with fixed $\lambda = 1.56$. The left panel is for the non-randomized setting and right panel corresponds to the randomized setting. Vertical bars are the barrier approximate credible intervals based on samples from the posterior using a Langevin random walk. Selection-adjusted point estimates – approximate MLE and approximate posterior mean – are marked on the bars.

6 Point estimates

In this section we explore point estimates for functions of the parameters β^* in the generative model. We can obtain an approximation for the posterior mean estimate by sampling from the approximate posterior as suggested in the previous section. We do implement this estimate when comparing different methods in Example 2 above, but the focus in this section is on studying the *mode* of the (approximate) selection-adjusted posterior. We refer to the latter as the selective MAP (or just the MAP) estimate. A main reason for concentrating on this estimate is that, by addressing the MAP, we make the discussion applicable also to the selective MLE (maximum-likelihood estimator), by considering a constant prior. We would like to point out that frequentist point estimation based on the full truncated likelihood (in particular, the MLE of β^*) is not addressed fully in the recent work on exact post-model selection inference, though there has been some work on pseudo-maximum likelihood estimators in Reid et al. (2014), as well as earlier work on two stage estimators Cohen and Sackrowitz (1989); Zhong and Prentice (2008); Zöllner and Pritchard (2007). Whereas implementation of the posterior mean estimate requires sampling from the approximate posterior, we will show that the problem of finding the approximate MAP is a convex optimization problem; this further adds to the appeal of the barrier approximation.

Recall that we proposed to approximate the selection-adjusted posterior in Section 4 by

$$\pi_S(\beta^*|y) \propto \pi(\beta^*) \frac{\exp\{-\|y - X^*\beta^*\|^2/(2\sigma^2)\}}{\exp(-h_\sigma(X^*\beta^*))} \quad (21)$$

where the function $h_\sigma(\mu)$ depends on A_E, b_E (and X^*, σ^2) and has the general form

$$h_\sigma(\mu) = \inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2\sigma^2} \|\mu - z\|^2 + w_\sigma(z) \right\},$$

and where we considered two specific choices for $w_\sigma(z)$ in (15) and (17). Fixing a choice of $w_\sigma(z)$, we will refer to a maximizer of (21) or, equivalently, a minimizer of

$$\left(\frac{1}{2\sigma^2} \|y - X^*\beta^*\|^2 - h_\sigma(X^*\beta^*) \right) - \log \pi(\beta^*), \quad (22)$$

as an approximate MAP estimator.

In the remainder of this section we assume without loss of generality that $\sigma^2 = 1$, i.e., $Y \sim N_n(X^*\beta^*, I)$. The following lemma asserts that finding an approximate MAP amounts to minimizing a convex objective.

Lemma 6.1 (Convexity of MAP (MLE) problem with general approximation). *Let $\pi(\beta^*)$ be a log-concave prior. For a nonnegative function $w(z)$, finding the minimizer of (22) in β^* is a convex optimization problem.*

Proof. With $\sigma^2 = 1$, minimizing (22) is equivalent to minimizing

$$\frac{1}{2} \|X^*\beta^*\|^2 - h(X^*\beta^*) - \beta^{*T}(X^{*T}y) - \log \pi(\beta^*).$$

Now

$$\begin{aligned} \frac{1}{2} \|X^*\beta^*\|^2 - h(\beta^*) &= \frac{1}{2} \|X^*\beta^*\|^2 - \inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2} \|X^*\beta^* - z\|^2 + w(z) \right\} \\ &= \sup_{z \in \mathbb{R}^n} \frac{1}{2} \|X^*\beta^*\|^2 - \frac{1}{2} \|X^*\beta^* - z\|_2^2 - w(z) \\ &= \sup_{z \in \mathbb{R}^n} z^T(X^*\beta^*) - \frac{1}{2} \|z\|^2 - w(z) \\ &= \bar{w}^*(X^*\beta^*) \end{aligned}$$

where \bar{w}^* is the convex conjugate of the function of $\bar{w}(z) = w(z) + \frac{1}{2}\|z\|^2$. Hence a MAP estimate minimizes

$$\bar{w}^*(X^*\beta^*) - \beta^{*T}(X^{*T}y) - \log \pi(\beta^*) \quad (23)$$

which is convex. \square

Remark 3. We quickly note from the above calculations that the gradient of the logarithm of the approximate selection adjusted posterior with a prior π on β^* can be computed as

$$\begin{aligned}\nabla \log \pi_S(\beta^*|y) &= X^{*T}(y - \nabla \bar{w}^*(X^*\beta^*)) + \nabla \log \pi(\beta^*) \\ &= X^{*T}(y - z^*(X^*\beta^*)) + \nabla \log \pi(\beta^*)\end{aligned}\tag{24}$$

where $z^*(\mu) = \operatorname{argsup}_{z \in \mathbb{R}^n} z^T \mu - \frac{1}{2} \|z\|^2 - w(z)$, the optimizer of the approximation. We use the above gradient of the log-posterior when drawing samples using a Langevin sampler.

Next we show that, at least under the saturated model, the minimizer of (23) has an explicit form.

Theorem 6.2. Under $Y \sim N_n(\beta^*, I)$, the approximate selective MAP satisfies

$$\hat{\beta}^* = \nabla \bar{w}(y + \nabla \log \pi(\hat{\beta}^*))\tag{25}$$

with $\bar{w}(z) = w(z) + \frac{1}{2} \|z\|_2^2$. In particular, taking $\pi(\beta^*) \propto 1$, the selective MLE is given by

$$\hat{\beta}^* = y + \nabla w(y)$$

Proof. Substituting $X^* = I_n$ into (23) and taking derivative w.r.t. β^* , the selective MAP satisfies

$$-y + \nabla \bar{w}^*(\hat{\beta}^*) - \nabla \log \pi(\hat{\beta}^*) = 0$$

which implies that $\hat{\beta}^* = \nabla \bar{w}^{*-1}(y + \nabla \log \pi(\hat{\beta}^*)) = \nabla \bar{w}(y + \nabla \log \pi(\hat{\beta}^*))$ where

$$\nabla \bar{w}(z) = z + \nabla w(z).$$

Now the selective MLE is the selective MAP with $\pi(\mu) = 1$, in which case $\log \pi(\hat{\mu}) = 0$ and the above reduces to

$$\hat{\beta}^* = \nabla \bar{w}(y) = y + \nabla w(y).$$

□

To implement the MAP estimate we need to specify a function $w(z)$. The following example shows that choosing $w(z) = I_{\mathcal{K}}(z)$, the first suggestion in Section 4, produces estimates that lack an important feature that the exact MAP has. After that, we will see that this undesirable situation is mitigated by replacing the indicator function with $w^{\text{bar}}(z)$.

Example 3. $Y \sim N(\beta^*, 1)$. For the selection event $\{y : y > 0\}$, the negative log of the selection-adjusted likelihood is $\frac{1}{2}(y - \beta^*)^2 + \log(1 - \Phi(-\beta^*))$ and the exact selective MLE can be easily seen to satisfy

$$y = \hat{\beta}^* + \frac{\phi(\hat{\beta}^*)}{1 - \Phi(-\hat{\beta}^*)}.$$

The selective MLE exhibits shrinkage of the unadjusted MLE, y , for y values close to the selection boundary; see the solid black line in the right panel of Figure 1. For any choice of $w(z)$, the approximate selective MLE solves

$$\underset{\beta^*}{\text{minimize}} \frac{1}{2} (\beta^{*2} - \mu y) - h(\beta^*) \quad (26)$$

with $h(\beta^*) = \inf_{z \in \mathbb{R}} \frac{1}{2} (\beta^* - z)^2 + w(z)$. If $w(z) = I_{\mathcal{K}}(z)$, then by Theorem 6.2, the approximate selective MLE is

$$\hat{\beta}^* = y + \nabla I_{\mathcal{K}}(y) = y,$$

noting that $y \in \mathcal{K}$. Hence, for that choice of $w(z)$ the approximate MLE coincides with the unadjusted MLE (In fact, the preceding argument can be extended to the general regression situation, observing that the minimizer of the approximate selective likelihood restricted to $\{\beta^* : X^* \beta^* \in \mathcal{K}\}$, is the usual least squares estimate $\hat{\beta}_{LS}^* = (X^{*T} X^*)^{-1} X^{*T} y$, whenever $X^* \hat{\beta}_{LS}^* \in \text{int}(\mathcal{K})$).

The barrier function defined in (17) has the property that bigger penalty is applied to values $z \in \mathcal{K}$ closer to the selection boundary. In our univariate case ($\sigma = 1$) the barrier function is just $w^{\text{bar}}(z) = \log(1 + 1/z)$ for $z > 0$ ($w^{\text{bar}}(z) = \infty$ otherwise). By Theorem 6.2, the corresponding approximate selective MAP is ($y > 0$)

$$\hat{\beta}^* = y + \nabla w_{\text{bar}}(y) = y - \frac{1}{y(y+1)}, \quad (27)$$

which indeed shrinks the unadjusted MLE towards zero. The broken line in the right panel of Figure 1 shows the estimator (27), obtained with the barrier function. We can see that it follows the true selective MLE closely, shrinking to $-\infty$ near the selection boundary while reverting to the unadjusted MLE as the observation lies deeper inside the selection region.

Theorem 6.2 applies to the saturated model. In that case, the approximate selective MLE with $w(z) = w^{\text{bar}}(z)$ has a simple, closed form. For the more general linear model we offer the following lemma, which is proved in the appendix.

Lemma 6.3. Suppose that $Y \sim N_n(X^* \beta^*, I)$ and the selection event is of the form $\mathcal{K} = \{Ay \leq b\} = \{y : a_i^T y - b_i \leq 0, i = 1, \dots, m\}$. The selective MAP satisfies the estimating equation

$$X^{*T} \nabla \bar{w}^*(X^* \hat{\beta}^*) = X^{*T} z^*(X^* \hat{\beta}^*) = X^{*T} y + \nabla \log \pi(\hat{\beta}^*)$$

and the selective MLE satisfies

$$\nabla \bar{w}^*(X^* \hat{\beta}^*) = z^*(X^* \hat{\beta}^*) = y + v_0$$

for $v_0 \in \mathcal{N}(X^{*T})$, the null space of X^{*T} , with

$$z^*(\mu) = \nabla \bar{w}^{-1}(\mu) = \underset{z \in \mathbb{R}^n}{\text{argsup}} z^T \mu - \frac{1}{2} \|z\|^2 - \sum_{j=1}^m \log(1 + 1/(b_i - a_i^T z)).$$

6.1 Selective MAP (MLE) with randomization

The MAP estimator for the randomized situation with additive, independent gaussian noise $W \sim N_n(0, \gamma^2 I)$ maximizes (20) with respect to β^* . By the comment in the last paragraph of Section 5.2, the approximate MAP for any w minimizes (22) only with $h(X^* \beta^*)$ replaced by (remember that $\sigma^2 = 1$)

$$\tilde{h}(X^* \beta^*) = \inf_{z \in \mathbb{R}^n} \left\{ \frac{1}{2(1 + \gamma^2)} \|X^* \beta^* - z\|^2 + w_{(1+\gamma^2)^{1/2}}(z) \right\}.$$

We can look for a maximizer by numerically optimizing the objective, relying on numerical methods to evaluate $\tilde{h}(X^* \beta^*)$ for a candidate β^* . The optimization problem is no longer equivalent to (23) because the variances used in the (unadjusted) likelihood and in \tilde{h} are not the same. Still, with our approximation, finding a MAP (MLE) is again a convex problem when a log-concave prior is used. Indeed, the negative of the logarithm of the approximate selection adjusted likelihood is

$$\begin{aligned} & -yX^* \beta^* + \frac{\gamma^2}{2(1 + \gamma^2)} \|X^* \beta^*\|_2^2 + \sup_{z \in \mathbb{R}^n} z^T \frac{X^* \beta^*}{1 + \gamma^2} - \left\{ \frac{\|z\|^2}{2(1 + \gamma^2)} + w_{(1+\gamma^2)^{1/2}}(z) \right\} \\ & = -yX^* \beta^* + \frac{\gamma^2}{2(1 + \gamma^2)} \|X^* \beta^*\|_2^2 + \bar{w}^* \left(\frac{X^* \beta^*}{1 + \gamma^2} \right), \end{aligned}$$

with

$$\bar{w}(z) = \frac{\|z\|^2}{2(1 + \gamma^2)} + w_{(1+\gamma^2)^{1/2}}(z). \quad (28)$$

Under the saturated model we can characterize the approximate MAP for the randomized setup by an estimating equation in the following lemma. The proof is seen trivially by taking the derivative of the logarithm of the selection-adjusted posterior.

Lemma 6.4 (Approximate selective randomized MAP under the saturated model). *Suppose that $Y \sim N_n(\beta^*, I)$ and let $\pi(\beta^*)$ be a log-concave prior. Let $\hat{\beta}^*$ be a minimizer of*

$$\frac{1}{2} \|y - \beta^*\|^2 - \tilde{h}(\beta^*) - \log \pi(\beta^*)$$

with \tilde{h} defined above. Then

$$\hat{\beta}^* = (1 + \gamma^2) \nabla \bar{w} \left((1 + \gamma^2)(y + \nabla \log \pi(\hat{\beta}^*)) - \gamma^2 \hat{\beta}^* \right)$$

with \bar{w} given by (28). In particular, the estimating equation for an approximate selective MLE is obtained by setting $\nabla \log \pi(\hat{\mu}) = 0$ above.

Example 3 (continued). $Y \sim (\beta^*, 1)$ and $W \sim N(0, \gamma^2)$, independent. Also $\mu \sim \pi$ where π is log-concave. The observation is Y truncated to $\{Y + W > 0\}$. We consider the approximate selective MLE, $\hat{\beta}^*$, corresponding to

$$w_{\sqrt{1+\gamma^2}}^{\text{bar}}(z) = \begin{cases} \log \left(1 + \frac{\sqrt{1+\gamma^2}}{z} \right) & z \in \mathcal{K} \\ \infty & \text{otherwise} \end{cases}.$$

By Lemma 6.4, this is given by

$$\hat{\beta}^* = (1 + \gamma^2) \left(\frac{z}{1 + \gamma^2} - \frac{\sqrt{1 + \gamma^2}}{z(z + \sqrt{1 + \gamma^2})} \right) \Big|_{z=Z_{obs}} \quad \text{with } Z_{obs} = (1 + \gamma^2)Y - \gamma^2\hat{\beta}^*.$$

The right panel of Figure 1 shows the randomized approximate MLE versus the approximate MLE (and other estimators) without randomization. Note that with randomization, y takes on values between $\pm\infty$ on the selection event (because y^* , not y , is restricted to be positive). Since the effect of truncation is attenuated by randomization, leaving more information for inference, shrinkage near the selective boundary is much more moderate than in the nonrandomized version.

Similarly to the relation between Theorem 6.2 and Lemma 6.3, there is an easy extension of 6.4 to the more general linear regression scenario in the randomized setting. The more general estimating equation for the randomized MAP $\hat{\beta}^*$ when $Y \sim N(X^*\beta^*, I_n)$ and $Y^* = Y + W$ where $W \sim N_n(0, \gamma^2 I_n)$ independent of Y can be written as

$$\frac{1}{(1 + \gamma^2)} X^{*T} \nabla w^* \left(\frac{X^* \hat{\beta}^*}{(1 + \gamma^2)} \right) = X^{*T} y + \nabla \log \pi(\hat{\beta}^*) - \frac{\gamma^2}{1 + \gamma^2} X^{*T} X^* \hat{\beta}^*$$

where

$$\nabla \bar{w}^*(\mu) = z^*(\mu) = \operatorname{argsup}_{z \in \mathbb{R}^n} z^T \mu - \frac{\|z\|^2}{2(1 + \gamma^2)} - w_{(1+\gamma^2)^{1/2}}(z).$$

7 Asymptotic guarantees under approximation

In this section we study asymptotic properties related to the barrier approximation to the adjustment factor, in an attempt to provide some theoretical support for using it. Specifically, we show first that for i.i.d. variables (from an arbitrary distribution) the barrier approximation to the adjustment factor has the correct limiting exponential decay rate under so-called non-local alternatives. The second main result of this section states that the randomized selective MLE of Section 6.1, which employs the barrier approximation, is a consistent estimator for the mean under non-local alternatives. Other common estimators lack consistency properties under the nonrandomized selective law; the unadjusted sample mean is trivially inconsistent, and we show that the selective MLE is also inconsistent for non-local alternatives. This again highlights the potential advantages in using randomization in selection.

The first result we state is a standard result in large deviations theory. For $d = 1$ the theorem below is known as Cramer's theorem, see Dembo and Zeitouni (2010). The result for higher dimensions uses similar arguments; except that the upper bound in that case requires a minimax result to interchange the supremum and infimum. Although there is nothing novel in Theorem 7.1, we provide a proof in the Appendix, which is adapted from Dembo and Zeitouni (2010).

Theorem 7.1. Let $Y_1, Y_2, \dots, Y_n \stackrel{i.i.d.}{\sim} F$ with support on \mathbb{R}^d , d fixed such that $F \in \mathcal{F} = \{F : \mathbb{E}_F(Y) = \beta^* \in \mathbb{R}^d, \mathbb{E}_F(Y - \beta^*)^2 = I_d\}$. For any convex set $\mathcal{K} \subseteq \mathbb{R}$ when $d = 1$, and for any convex and compact set $\mathcal{K} \subseteq \mathbb{R}^d$ when $d > 1$, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) = - \inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*) \quad (29)$$

where

$$\Lambda^*(x; \beta^*) = \sup_{\eta} \eta^T x - \log \mathbb{E}[(\exp(\eta^T Y_1) | \beta^*)]$$

is the convex conjugate of the log of the moment generating function $\Lambda(\eta; \beta^*) = \log \mathbb{E}[\exp(\eta^T Y_1) | \beta^*]$ of random variable Y_1 .

As a consequence we have the following corollary that says that the sequence of selection-adjusted posterior based on a suitable barrier approximation indeed converges to the corresponding (true) selection-adjusted posterior.

Corollary 7.2. Let $Y_1, Y_2, \dots, Y_n \sim N_d(\beta^*, I)$ be random d -dimensional Normal vectors. Consider the selection event $\{\bar{Y}_n \in \mathcal{K}\}$ for a convex (and compact for $d > 1$) set $\mathcal{K} \subseteq \mathbb{R}^d$. Denote by

$$\log \pi_S(\beta^* | \bar{y}_n) = \log \pi(\beta^*) - n \|\bar{y}_n - \beta^*\|^2 / 2 - \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*)$$

the corresponding (true) selection-adjusted posterior, and by

$$\log \hat{\pi}_S(\beta^* | \bar{y}_n) = \log \pi(\beta^*) - n \|\bar{y}_n - \beta^*\|^2 / 2 + n \inf_{x \in \mathbb{R}^d} \left\{ \|x - \beta^*\|^2 / 2 + \frac{1}{n} w_{1/\sqrt{n}}(x) \right\}$$

a corresponding approximate selection-adjusted posterior, where $w_{1/\sqrt{n}}(\cdot)$ is a (convex) barrier function associated with \mathcal{K} . Then

$$\lim_{n \rightarrow \infty} \frac{1}{n} \{\log \pi_S(\beta^* | \bar{y}_n) - \log \hat{\pi}_S(\beta^* | \bar{y}_n)\} \rightarrow 0 \text{ as } n \rightarrow \infty$$

for all $\beta^* \in \mathbb{R}^d$ whenever $\frac{1}{n} w_{1/\sqrt{n}}(\cdot)$ converges to $I_{\mathcal{K}}(\cdot)$ pointwise as $n \rightarrow \infty$.

Remark 4. The results in this section extend to the case of general regression with random design,

$$(x_i, y_i) \stackrel{i.i.d.}{\sim} F, (x_i, y_i) \in \mathbb{R}^p \times \mathbb{R}$$

with a fixed p . In that case, the selection probability typically takes on the form

$$\mathbb{P}(A_E \sqrt{n} \bar{T}_n \geq b_E) = \mathbb{P}(A_E \sqrt{n} (\bar{T}_n - \beta^*) \geq b_E - \sqrt{n} A_E \beta^*)$$

for a statistic $\bar{T}_n = \bar{T}_n(X, Y) \in \mathbb{R}^p$; a polytope defined by $A_E \in \mathbb{R}^{m \times p}, b_E \in \mathbb{R}^m$; and a population parameter $\beta^* = \mathbb{E}(\bar{T}_n)$. The statistic $\bar{T}_n(X, Y)$ is usually asymptotically Gaussian and the central limit theorem can be again used when $\sqrt{n} \beta^* = O(1)$ to obtain the selection probability as $n \rightarrow \infty$. An example is best linear prediction, where the object of inference is

$$\beta_E^* = \mathbb{E}_F(X_E^T X_E)^{-1} \mathbb{E}_F(X_E^T Y),$$

where Lasso selects the active set of coefficients E . The event of selecting set E with corresponding signs z_E can be described as an affine event of the above form with

$$\bar{T}_n = \begin{pmatrix} \hat{\beta}_E \\ \frac{1}{n} X_{-E}^T (y - X_E \hat{\beta}_E) \end{pmatrix}$$

for $\hat{\beta}_E = (X_E^T X_E)^{-1} X_E^T Y$. The statistic \bar{T}_n is asymptotically Gaussian with mean

$$\begin{pmatrix} \beta_E^* \\ \mathbb{E}_F(X_{i,-E}^T (y_i - X_{i,E} \beta_E^*)) \end{pmatrix}$$

and a finite covariance matrix Σ/n . A similar example is discussed in the randomized setup in [Tian and Taylor \(2015\)](#) with the logistic loss. For non-local alternatives, namely, when $\|b_E - \sqrt{n} A_E \beta^*\| \rightarrow \infty$ as $n \rightarrow \infty$, we are back to the univariate problem illustrated above. The proposed approximation gives that

$$\log \mathbb{P}(\bar{T}_n \in \mathcal{K} | \beta^*) \sim -n \inf_{x \in \mathcal{K}} \left\{ \Lambda^*(x, \beta^*) + \frac{1}{n} w_{1/\sqrt{n}}(x) \right\}$$

for

$$\Lambda^*(x; \beta^*) = \sup_{\eta} \eta^T x - \log \mathbb{E}[(\exp(\eta^T T) | \beta^*)] \text{ and } \mathcal{K} = \{z : A_E z \geq 0\}.$$

In terms of notation, it would make sense to switch to \bar{T}_n to denote the sample mean of an asymptotically Gaussian statistic, but we state all results for $\bar{Y}_n \in \mathbb{R}^d$ in order to avoid introducing additional notation.

We now shift the focus to the *randomized* approximate selective MLE discussed in Section 6.1 and show that it is a consistent estimator for β^* under non-local alternatives. Before proceeding, it is important to note that the *nonrandomized* approximate selective MLE based on the barrier approximation (17) is generally not consistent for non-local alternatives:

Theorem 7.3. *Let $Y_i \sim N(\beta^*, 1)$ be i.i.d. random variables and consider inference for β^* conditionally on $\{\sqrt{n} \bar{Y}_n > 0\}$. Then the maximizer $\hat{\beta}_n^*$ of the approximate selection-adjusted log likelihood,*

$$\ell_S^n(\beta^*) = -n(\bar{y}_n - \beta^*)^2 + n \inf_{x \in \mathbb{R}} \left\{ \frac{(x - \beta^*)^2}{2} + \frac{1}{n} \log \left(1 + \frac{1}{\sqrt{n} \bar{y}_n} \right) \right\}, \quad (30)$$

does not converge in probability to β^ when $\beta^* = -\delta < 0$. More generally, $\hat{\beta}_n^*$ is inconsistent for the sequence of parameters β_n^* under the asymptotics $\sqrt{n} \beta_n^* \rightarrow -\infty$.*

Remark 5. *Theorem 7.3 is stated for the barrier approximation (17) for clarity of exposition. The approximate selective MLE is not consistent as long as the function $w(\cdot)$ satisfies*

$$x \nabla w(x) \rightarrow C \text{ as } x \downarrow 0.$$

A remark is also in order regarding the unadjusted estimator \bar{Y}_n . [Tian and Taylor \(2015\)](#) showed that there are heavier tailed randomization schemes under which \bar{Y}_n is a consistent estimator for non-local alternatives. We note that this is not true for Gaussian randomization. Indeed, let $\beta^* < 0$ be fixed (with n) and the selection event be $\{\sqrt{n}\bar{Y}_n + W > 0\}$, where $W \sim N(0, \gamma^2)$ independently of \bar{Y}_n . Then it can be shown that, for large n , the selective distribution of $\sqrt{n}(\bar{Y}_n - \beta^*)$ is approximately $\mathcal{N}(-\sqrt{n}\beta^*/(1 + \gamma^2), (1 + 1/\gamma^2)^{-1})$, and hence \bar{Y}_n cannot be consistent for β^* .

Denote by

$$\ell_S^n(\beta^*) = \log f_S^n(\bar{y}_n|\beta^*) = -n\|\bar{y}_n - \beta^*\|^2/2 - \log \mathbb{P}(\bar{Y}_n^* \in \mathcal{K}|\beta^*), \quad (31)$$

the (exact) selection-adjusted posterior that corresponds to applying selection to $\bar{Y}_n^* \sim \mathcal{N}(\beta^*, 1 + \gamma^2)$, the randomized sample mean. Denote also by

$$\hat{\ell}_S^n(\beta^*) = \log \hat{f}_S^n(\bar{y}_n|\beta^*) = -n\|\bar{y}_n - \beta^*\|^2/2 + n \inf_{x \in \mathbb{R}^d} \left\{ \frac{\|x - \beta^*\|^2}{2(1 + \gamma^2)} + \frac{1}{n} w \sqrt{1 + \gamma^2}/\sqrt{n}(x) \right\} \quad (32)$$

the corresponding approximate selection-adjusted posterior. We make a few crucial observations about the approximate sequence of likelihoods $\ell_S^n(\beta^*)$ before stating the main result.

Lemma 7.4 (Strong convexity). *The sequence of approximate log partition functions*

$$\hat{C}_n(\beta^*) = \frac{\gamma^2 \|\beta^*\|^2}{2(1 + \gamma^2)} + \bar{w}_n^*(\beta^*) \quad (33)$$

where $\bar{w}_n^*(\cdot)$ is the convex conjugate of

$$\bar{w}_n(x) = \|x\|^2/2(1 + \gamma^2) + \frac{1}{n} w \sqrt{1 + \gamma^2}/\sqrt{n}(x)$$

and hence, the negative of logarithms of approximate selection adjusted likelihood $\hat{f}_S^n(\beta^*)$ are strongly convex with indices of convexity m_n lower bounded by $M = \gamma^2/(1 + \gamma^2)$.

Remark 6 (Uniform convergence on compact sets). *The approximation $\hat{\ell}_S^n(\beta^*)$ is continuous in β^* and so is its limit. Hence the sequence of scaled differences $n^{-1}(\hat{\ell}_S^n(\beta^*) - \ell_S^n(\beta^*))$ converges uniformly on a compact subset Θ of the parameter space.*

Lemma 7.5. *If \hat{C}_n is the approximate log partition function in (33), then the approximate randomized selective MLE $\hat{\beta}_n^*$ that maximizes (32) satisfies the following inequality*

$$\|\hat{\beta}_n^* - \beta^*\|^2 \leq \frac{(\gamma^2 + 1)^2}{\gamma^4} \|\bar{y}_n - \nabla \hat{C}_n(\beta^*)\|^2 \text{ for all } \beta^* \in \mathbb{R}^p.$$

We are now ready to prove consistency of the approximate randomized selective MLE.

Theorem 7.6. *For a convex (and compact for $d > 1$) set $\mathcal{K} \subset \mathbb{R}^d$, the approximate randomized selective MLE β_n^* based on maximizing $\hat{\ell}_S^n(\beta^*)$ in (32) is consistent for parameter β^* under the selective law (31).*

To complete the picture, we prove that the approximate randomized selective MLE converges to the exact randomized selective MLE.

Theorem 7.7. *If $\hat{\beta}_n^*$ denotes the approximate randomized selective MLE, obtained as the maximizer of $\hat{\ell}_S^n(\beta^*)$ in (32) and $\tilde{\beta}_n^*$ is the selective MLE, which maximizes $\ell_S^n(\beta^*)$ in (31), then for any $\delta > 0$,*

$$\mathbb{P}(\|\hat{\beta}_n^* - \tilde{\beta}_n^*\| > \delta \mid \bar{Y}_n^* \in \mathcal{K}) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

8 Discussion

In the current work we adopted the perspective in which post-model selection inference is modeled as a truncation problem. Hence, inference regarding any specific parameter should be valid under the conditional distribution of the data given that this parameter was selected. When model selection is done by a known formal procedure specified before observing the data, the conditional approach aims for the exact distribution of the data; therefore, it may lead to more precise inference as compared to methods which offer guarantees irrespective of the selection procedure (for example, Berk et al., 2013; Dwork et al., 2015); the latter have their own advantages, especially the fact that they do not depend on the selection algorithm.

A recent line of work (Lee et al., 2013; Taylor et al., 2013, 2014; Lee and Taylor, 2014, for example) has contributed to the feasibility of conditional inference after variable selection in generalized linear models; these contributions are essentially based on the observation that classical theory for exponential families applies also to the truncated distribution (Fithian et al., 2014). However, it is mainly testing (and confidence intervals) that is addressed by this theory and pursued in recent works; other questions remain open. For example, this theory has little implication on how to construct “good” point estimators for the regression coefficients after model selection.

Since, by conditioning, the starting point of the inference stage is *after* selection, it makes as much sense to use a prior for the parameters of the selected model as it does in the non-adaptive situation. The Bayesian approach has the advantage that any kind of inference (estimation or hypothesis testing, for example) is based on the posterior distribution of the parameters given the data: in that sense, Bayesian inference is more straightforward. We would like to summarize and point out some further advantages of selection-adjusted Bayesian approach:

1. Even after substituting our approximation (or any other) to the adjustment factor, frequentist inference based on the approximate truncated likelihood is challenging, to say the least. The conditions for usual asymptotic theory (say, for the likelihood ratio statistic) to hold need to be checked carefully. On the other hand, when taking

a Bayesian approach we can resort to posterior sampling techniques without relying on asymptotics. As a concrete example consider constructing confidence sets for two coefficients or more after selection; the implementation of UMPU tests in [Fithian et al. \(2014\)](#) is considerably more difficult because the problem now does not reduce to constructing a CI for a univariate truncated normal. By contrast, extending the procedures in [Section 5](#) to construct credible sets instead of intervals is straightforward.

2. With the Bayesian approach we can easily obtain approximate posterior predictive intervals under the selected model. Hence, if $\tilde{Y}|\beta^* \sim N(\tilde{x}^T\beta^*, \sigma^2)$ is independent of $Y \in \mathbb{R}^n$, then

$$f_S(\tilde{y}|y) = \int f_S(\tilde{y}, \beta^*|y) d\beta^* = \int f(\tilde{y}|\beta^*) \pi_S(\beta^*|y) d\beta^*$$

(the unadjusted likelihood is used for \tilde{y}). As usual, this is the marginal distribution of \tilde{Y} with respect to the updated distribution of the parameter of the model, only that observed data Y updates the prior via the *selection-adjusted* posterior. Sampling from the predictive distribution is trivial after obtaining (per the prescription in [Section 5](#), for example) a sample from the approximate posterior distribution .

3. The often inconvenient assumption of known noise level can be replaced with a prior distribution on σ^2 . Our methods can be extended to apply to generalized linear models, and hence can in principle accommodate the situation where there is a prior on σ^2 . A non-informative prior may be preferred in this situation.

A main message which we tried to convey, then, is that methods suggested recently for conditional inference in the linear model are not all-encompassing; and when existing tools fail to provide conclusive answers, the Bayesian approach pursued in the current paper can be useful.

There are certainly directions for further research. It would be interesting to study the performance of the approximate point estimators of [section 6](#) in finite samples. An important question is what benchmark can be used for the mean squared error of an estimator. From a frequentist perspective, there is usually no unbiased estimator, and the Cramer-Rao bound is not applicable (even if we knew how to compute it) in the small-sample regime. From a Bayesian viewpoint, the minimum Bayes risk – the expectation of the conditional variance – usually cannot be computed analytically either, but tractable lower bounds on it (which would automatically bound the Bayes risk of any other estimator) may exist. Lastly, other suggestions for an approximation to the truncated likelihood may be considered, and it would be interesting to see how sensitive the adjusted posterior is to that choice.

Acknowledgement. Providing Bayesian adjusted inference after variable selection was previously proposed by Daniel Yekutieli and Edward George and presented at the 2012 Joint Statistical Meetings in San Diego. To a large extent, the framework we proposed in [Section 2](#) relies on their ideas, which we greatly acknowledge. Our interest re-arose with the recent developments on exact post-selection inference in the linear model. We are thankful

to Daniel and Ed for helpful conversations and to Ed for pointing out the reference to [Bayarri and DeGroot \(1987\)](#). We would also like to sincerely thank and acknowledge Chiara Sabatti for the long discussions and for her suggestions along the way, which greatly improved presentation helped to form our perspective on the problem of Bayesian post-model selection inference.

Supplementary Material. The supplement contains an analysis of real data and a section on using our methods in frequentist inference. Specifically, we utilize our methods to derive an approximate uniformly minimum variance unbiased estimator for the saturated model; and describe how our methods can be applied to approximate a uniformly most powerful unbiased test when the computation does not reduce to constructing tests for a univariate truncated normal variable.

9 Supplement

9.1 A real data example

We give Bayesian post-selection inference to the publicly available diabetes data set from [Efron et al. \(2004\)](#). The data set consists of 10 standardized measurements from 442 samples. We apply the Lasso at a fixed value of the tuning parameter, $\lambda = \hat{\sigma} \mathbb{E}(\|X^T \epsilon\|_\infty)$ where $\epsilon \sim N_n(0, I)$ (as proposed in [Negahban et al., 2009](#)), to select a subset of predictors. [Lee and Taylor \(2014\)](#) have also analyzed this data set and constructed confidence intervals based on the univariate truncated Gaussian pivot in the *saturated* model.

For the above choice of tuning parameter, the Lasso selects the set of variables $E = \{\text{sex, bmi, bp, s1, s3, s5, s6}\}$. We run selection again, first with randomized Gaussian response, adding an independent variable $W \sim N(0, \gamma^2 I)$ to the response; and second with a random 80% split of the data to construct data-carving intervals. We obtain credible intervals and point estimates – selective posterior mean and MAP – for $\beta^E = \mathbb{E}[(X_E^T X_E)^{-1} X_E^T Y]$ based on a non-informative prior and the *selected* model as the generative model. For comparison, we also plot the intervals of [Lee and Taylor \(2014\)](#), even though the models for inference are different. The randomized and data-carving intervals are comparable in length to the unadjusted ones, and much shorter than the proposed Bayesian adjusted intervals and the frequentist confidence intervals of [Lee and Taylor \(2014\)](#). Another advantage of the proposed adjusted intervals (randomized or not), is that, utilizing posterior sampling, we are able construct them for any generative model. By contrast, for the confidence intervals of [Lee and Taylor \(2014\)](#) it is important to assume the saturated model, because in that case a simple univariate pivot exists.




Figure 3: Selection-adjusted inference after selection via Lasso: the left panel gives Bayesian adjusted intervals for $\beta^E = \mathbb{E}[(X_E^T X_E)^{-1} X_E^T Y]$ under the selected model and a non-informative prior. Vertical bars are the Barrier approximate adjusted credible intervals from Section 5 of the main paper. The right panel shows the point estimates: MLE and mean based on (the approximate) selection adjusted posterior are marked on the bars.

9.2 Using our approximation in frequentist inference

Since our approximation serves as a surrogate to the selection-adjusted *likelihood*, we can use it to provide approximate inference from an entirely frequentist approach; as we did in Section 7 when considering the approximate maximum-likelihood (ML) estimate. In this section we discuss a few other cases where our barrier approximation is advantageous in frequentist computations. We consider frequentist procedures suggested in [Tian and Taylor \(2015\)](#) and in [Fithian et al. \(2014\)](#) that involve approximation of intractable integrals; we demonstrate how our methods can be used instead of generating random samples to approximate these integrals, and thereby reduce computational effort. In fact, our approximation performs well in recovering the densities against which these integrals are taken; hence, the density produced using our approximation can also serve as an effective choice for a proposal function when implementing MCMC techniques.

9.2.1 Selective UMVUE under randomized response

Consider inference under the saturated model, $Y \sim N_n(\beta^*, \sigma^2 I)$. In general, an unbiased estimator for β^* conditional on $Y \in \{y : A_E y \leq b_E\}$ does not exist. However, [Tian and Taylor \(2015\)](#) show that if selection is based on a randomized response, a simple unbiased estimator exists, which can be Rao-Blackwellized to produce an Unbiased Minimum-Variance Estimator (UMVUE) as suggested in [Fithian et al. \(2014\)](#). Indeed, let $Y^* = Y + W$

where $W \sim N_n(0, \gamma^2 I)$ is an independent randomization term, and let the selection event be $\{(y, w) : A_E(y + w) \leq b_E\}$. It is trivial to verify that, marginally,

$$\hat{\beta}^* = Y - \frac{\sigma^2}{\gamma^2} W$$

is independent of $Y^* = Y + W$ and unbiased for β^* . Therefore, it is unbiased for β^* also conditionally on $A_E Y^* \leq b_E$. Y is of course a sufficient statistic for β^* conditionally on selection, hence, invoking a Rao-Blackwell argument, the selective UMVUE is $T(Y)$ where

$$T(y) = \mathbb{E} \left(Y - \frac{\sigma^2}{\gamma^2} W \middle| y, A_E(Y + W) \leq b_E \right) = y - \frac{\sigma^2}{\gamma^2} \mathbb{E}(W | A_E(y + W) \leq b_E). \quad (34)$$

Thus, Rao-Blackwellization reduces to computing $\mathbb{E}(W | A_E(y + W) \leq b_E)$, which [Tian and Taylor \(2015, Section 3.2\)](#) suggest to numerically approximate by a Monte-Carlo estimate. Specifically, Hit-and-Run Gibbs sampling can be used to simulate draws from the conditional distribution of W truncated to $A_E W \leq b_E - A_E y$. Using our approximation to the adjustment factor we can offer an alternative approximation to the UMVUE without sampling, as shown by the following theorem.

Theorem 9.1 (Approximate selective UMVUE in saturated model). *Suppose that $Y \sim N(\beta^*, \sigma^2 I_n)$ and let $Y^* = Y + W$ where $W \sim N_n(0, \gamma^2 I)$ independent of Y . Then the approximate selective UMVUE for μ conditionally on $A_E Y^* \leq b_E$ is given by*

$$\tilde{T}(y) = \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} z^*(y)$$

where

$$z^*(y) = \arg \max_z \left\{ \frac{z^T y}{\gamma^2} - \left(\frac{\|z\|^2}{2\gamma^2} + \sum_{i=1}^m \log \left(1 + \frac{\gamma}{(b_i - a_i^T z)} \right) \right) \right\}. \quad (35)$$

Proof. We have

$$\begin{aligned} \mathbb{E} \left(Y - \frac{\sigma^2}{\gamma^2} W \middle| y, A_E(W + y) \leq b_E \right) &= y - \frac{\sigma^2}{\gamma^2} \mathbb{E}(W | y, A_E(y + W) \leq b_E) \\ &= \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} \mathbb{E}(W + y | y, A_E(W + y) \leq b_E) \\ &= \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} \mathbb{E}(R | y, A_E R \leq b_E) \end{aligned}$$

where $R = (W + y)$. The task is now to approximate the conditional expectation in the last equation. Since $R \sim N_n(\alpha, \gamma^2 I_n)$ with $\alpha = y$, the conditional distribution of R given $A_E R \leq b_E$, belongs to an exponential family with density

$$f(r | A_E R \leq b_E) \propto \exp \left(-\|r - \alpha\|^2 / 2\gamma^2 \right) I(A_E r \leq b_E).$$

The corresponding log-partition function is

$$A(\eta) = -\gamma^2 \frac{\eta^T \eta}{2} - \log \mathbb{P}(A_E R \leq b_E | \eta),$$

where $\eta = \alpha/\gamma^2$ is a re-parametrization of the natural parameter. With the barrier approximation for negative of $\log \mathbb{P}(A_E R \leq b_E)$ plugged in, we approximate the log-partition function $A(\eta)$ as

$$\sup_{z \in \mathbb{R}^n} z^T \eta - \left(\frac{\|z\|^2}{2\gamma^2} + \sum_{i=1}^m \log(1 + \gamma/(b_i - a_i^T z)) \right) = g^*(\eta),$$

the convex conjugate of $g(z) = \frac{\|z\|^2}{2\gamma^2} + \sum_{i=1}^m \log \left(1 + \frac{\gamma}{(b_i - a_i^T z)} \right)$. Then

$$\begin{aligned} \mathbb{E} \left(y - \frac{\sigma^2}{\gamma^2} W \middle| y, A_E(W + y) \leq b_E \right) &\approx \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} \frac{\partial A(\eta)}{\partial \eta} \bigg|_{y/\gamma^2} \\ &= \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} \nabla g^*(\eta) \big|_{y/\gamma^2} \\ &= \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} (\nabla g)^{-1}(\eta) \big|_{y/\gamma^2} \\ &= \left(1 + \frac{\sigma^2}{\gamma^2} \right) y - \frac{\sigma^2}{\gamma^2} z^*(y). \end{aligned}$$

The final equation follows from observing that

$$z^*(y) = \arg \max_z \left\{ \frac{z^T y}{\gamma^2} - \left(\frac{\|z\|^2}{2\gamma^2} + \sum_{i=1}^m \log \left(1 + \frac{\gamma}{(b_i - a_i^T z)} \right) \right) \right\}$$

satisfies

$$y/\gamma^2 = \nabla g(z^*(y)),$$

which implies

$$z^*(y) = (\nabla g)^{-1}(y/\gamma^2) = \nabla g^*(y/\gamma^2).$$

□

Example 5 (continued). *The setup is the randomized setup of Example 2 of Subsection 7.1. By Theorem 9.1, the approximate selective UMVUE is*

$$\begin{aligned} \hat{\beta}^* &\approx \left(1 + \frac{1}{\gamma^2} \right) y - \frac{1}{\gamma^2} z^*(y), \text{ where} \\ z^*(y) &= \arg \max_z \left\{ \frac{z^T y}{\gamma^2} - \left(\frac{\|z\|^2}{2\gamma^2} + \log(1 + \gamma/z) \right) \right\}. \end{aligned}$$

Figure 4 shows how closely the approximate UMVUE mimics the exact UMVUE in the univariate example.

The statement in Theorem 9.1 is for the saturated model.




Figure 4: Approximate UMVUE for the univariate example with randomization, $Y \sim N(\beta^*, 1)$ truncated to $\{Y + W > 0\}$.

9.2.2 Approximate UMPU tests under the selected model

Consider inference for β_j^* under a generative model, $Y \sim N(X^* \beta^*, \sigma^2 I)$, and conditionally on $\{\hat{E}(Y) = E\}$, equivalently, $\{A_E Y \leq b_E\}$. To present this section without notational challenges, let $X^* = X_{\tilde{E}}, \beta^* = \beta_{\tilde{E}}$ for $\tilde{E} \subset \{1, 2, \dots, p\}$. The special case of $\tilde{E} = E$ corresponding to $X^* = X_E$ leads to inference in the selected model. For known σ^2 , it follows from the general theory in [Fithian et al. \(2014\)](#) that Uniformly Most Powerful Unbiased (UMPU) tests (and confidence intervals) are based on the conditional law

$$\mathcal{L}(X_{j \cdot \tilde{E}}^T Y | A_E Y \leq b_E, X_{\tilde{E} \setminus j}^T Y) = \mathcal{L}(\eta^T Y | A_E Y \leq b_E, X_{\tilde{E} \setminus j}^T Y), \quad (36)$$

with

$$\eta = \frac{X_{j \cdot \tilde{E}}}{\|X_{j \cdot \tilde{E}}\|^2}, \quad X_{j \cdot \tilde{E}} = R_{\tilde{E} \setminus j} X_j$$

and $R_{\tilde{E} \setminus j} = (I - P_{\tilde{E} \setminus j})$ and $P_{\tilde{E} \setminus j} = X_{\tilde{E} \setminus j}^T (X_{\tilde{E} \setminus j}^T X_{\tilde{E} \setminus j})^{-1} X_{\tilde{E} \setminus j}^T$ is the projection matrix onto the column space of $X_{\tilde{E} \setminus j}$. In particular, this distribution depends only on $\beta_{j \cdot \tilde{E}}$, the j -th parameter in the model \tilde{E} . Denoting $T = \eta^T Y$, we observe that, unconditionally, $T \sim N(\beta_{j \cdot \tilde{E}}, \sigma^2 / \|X_{j \cdot \tilde{E}}\|^2)$. Finally, noting that

$$y = P_{\tilde{E} \setminus j} y + (t / \|\eta\|^2) \eta + P_{\tilde{E}}^\perp y \quad (37)$$

where $P_{\tilde{E}}^\perp = (I - P_{\tilde{E}})$, we conclude that

$$\mathcal{L}(\eta^T Y | A_E Y \leq b_E, X_{\tilde{E} \setminus j}^T y) \equiv \mathcal{L}(T | A_E \{(T / \|\eta\|^2) \eta + P_{\tilde{E}}^\perp Y\} \leq \tilde{b}, P_{\tilde{E} \setminus j} y)$$

with

$$\tilde{b} = b - AP_{\tilde{E} \setminus j} y.$$

This law belongs to a one parameter exponential family with density proportional to

$$\exp \left\{ -\|X_{j \cdot \tilde{E}}\|^2 (t - \beta_{j \cdot \tilde{E}})^2 / 2\sigma^2 \right\} \times \int_{\{v: A_E \{(t/\|\eta\|^2)\eta + v\} \leq \tilde{b}\}} \exp(-v^T \tilde{\Sigma}_V^{-1} v / 2) dv \quad (38)$$

where

$$V = (I - P_{\tilde{E}})Y \sim N_n(0, \Sigma_V), \quad \Sigma_V = \sigma^2(I - P_{\tilde{E}}).$$

As

$$\int_{\{v: A_E \{(t/\|\eta\|^2)\eta + v\} \leq \tilde{b}\}} \exp(-v^T \Sigma_V^{-1} v / 2) dv \propto \mathbb{P}(A_E V \leq b'(t))$$

where $b'(z) = \tilde{b} - (z/\|\eta\|^2)A_E \eta$, we can use our approximation to compute the integral. In this way will be able to compute the approximate density of T on a grid of values on the real line, and use this discretized version to obtain the approximate UMPU tests described in general in Theorem 5 of [Fithian et al. \(2014\)](#).

References

- MJ Bayarri and MH DeGroot. Bayesian analysis of selection models. *The Statistician*, pages 137–146, 1987.
- Yoav Benjamini and Amit Meir. Selective correlations-the conditional estimators. *arXiv preprint arXiv:1412.3242*, 2014.
- Richard Berk, Lawrence Brown, Andreas Buja, Kai Zhang, Linda Zhao, et al. Valid post-selection inference. *The Annals of Statistics*, 41(2):802–837, 2013.
- Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- Arthur Cohen and Harold B Sackrowitz. Two stage conditionally unbiased estimators of the selected mean. *Statistics & Probability Letters*, 8(3):273–278, 1989.
- Amir Dembo and Ofer Zeitouni. Large deviations techniques and applications, volume 38 of stochastic modelling and applied probability, 2010.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Preserving statistical validity in adaptive data analysis. *arXiv preprint arXiv:1411.2664*, 2014.
- Cynthia Dwork, Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Aaron Roth. Generalization in adaptive data analysis and holdout reuse. *arXiv preprint arXiv:1506.02629*, 2015.

- Bradley Efron. Tweedies formula and selection bias. *Journal of the American Statistical Association*, 106(496):1602–1614, 2011.
- Bradley Efron. *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press, 2012.
- Bradley Efron, Trevor Hastie, Iain Johnstone, Robert Tibshirani, et al. Least angle regression. *The Annals of statistics*, 32(2):407–499, 2004.
- William Fithian, Dennis Sun, and Jonathan Taylor. Optimal inference after model selection. *arXiv preprint arXiv:1410.2597*, 2014.
- E George and D Yekutieli. Shrinkage adjustment for model selection. Presented at the 2012 Joint Statistical Meetings, San-Diego, CA, 2012.
- Marek Kanter and Harold Proppe. Reduction of variance for gaussian densities via restriction to convex sets. *Journal of Multivariate Analysis*, 7(1):74–81, 1977.
- J Lee, D Sun, Yuekai Sun, and J Taylor. Exact post-selection inference with the lasso. *arXiv preprint arXiv:1311.6238*, 2013.
- Jason D Lee and Jonathan E Taylor. Exact post model selection inference for marginal screening. In *Advances in Neural Information Processing Systems*, pages 136–144, 2014.
- Joshua R Loftus. Selective inference after cross-validation. *arXiv preprint arXiv:1511.08866*, 2015.
- Joshua R Loftus and Jonathan E Taylor. Selective inference in regression models with groups of variables. *arXiv preprint arXiv:1511.01478*, 2015.
- Micha Mandel and Yosef Rinott. On statistical inference under selection bias. 2007.
- Sahand Negahban, Bin Yu, Martin J Wainwright, and Pradeep K Ravikumar. A unified framework for high-dimensional analysis of m -estimators with decomposable regularizers. In *Advances in Neural Information Processing Systems*, pages 1348–1356, 2009.
- Stephen Reid, Jonathan Taylor, and Robert Tibshirani. Post-selection point and interval estimation of signal sizes in gaussian samples. *arXiv preprint arXiv:1405.3340*, 2014.
- H Robbins. The uv method of estimation. *Statistical Decision Theory and Related Topics IV*, 1:265–270, 1988.
- Daniel Russo and James Zou. Controlling bias in adaptive data analysis using information theory. *arXiv preprint arXiv:1511.05219*, 2015.
- Noah Simon and Richard Simon. On estimating many means, selection bias, and the bootstrap. *arXiv preprint arXiv:1311.3709*, 2013.

- Uri Simonsohn, Leif D Nelson, and Joseph P Simmons. P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143(2):534, 2014.
- J Taylor, R Lockhart, RJ Tibshirani, and R Tibshirani. Exact post-selection inference for sequential regression procedures. *Journal of the American Statistical Association*, 111(514):600–620, 2014.
- Jonathan Taylor, Joshua Loftus, and Ryan Tibshirani. Tests in adaptive regression via the kac-rice formula. *arXiv preprint arXiv:1308.3020*, 2013.
- Xiaoying Tian and Jonathan E Taylor. Selective inference with a randomized response. *arXiv preprint arXiv:1507.06739*, 2015.
- Kenneth W Wachter and James Trussell. Estimating historical heights. *Journal of the American Statistical Association*, 77(378):279–293, 1982.
- Asaf Weinstein, William Fithian, and Yoav Benjamini. Selection adjusted confidence intervals with more power to determine the sign. *Journal of the American Statistical Association*, 108(501):165–176, 2013.
- Fan Yang, Rina Foygel Barber, Prateek Jain, and John Lafferty. Selective inference for group-sparse linear models. In *Advances in Neural Information Processing Systems*, pages 2469–2477, 2016.
- Daniel Yekutieli. Adjusted bayesian inference for selected parameters. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 74(3):515–541, 2012.
- Hua Zhong and Ross L Prentice. Bias-reduced estimators and confidence intervals for odds ratios in genome-wide association studies. *Biostatistics*, 9(4):621–634, 2008.
- Sebastian Zöllner and Jonathan K Pritchard. Overcoming the winners curse: estimating penetrance parameters from case-control data. *The American Journal of Human Genetics*, 80(4):605–615, 2007.

A Proofs

Proof of Theorem 4.1. In the theorem presented below, we obtain a Chernoff bound for a random variable Z from an exponential family supported on \mathbb{R}^d and with natural parameter $\beta^* \in \mathbb{R}^d$ and a log-MGF $\Lambda(\eta; \beta^*)$, evaluated at η . Theorem 4.1 will follow as a corollary.

Theorem A.1. *If $Z \sim \exp(\beta^{*T}Z - H(\beta^*))$, then for a convex, compact, selective region $\mathcal{K} \subset \mathbb{R}^d$, the selection probability of the event $\{z : z \in \mathcal{K}\}$ is bounded above as*

$$\log \mathbb{P}_{\beta^*}(Z \in \mathcal{K}) \leq - \inf_{z \in \mathcal{K}} \left\{ \sup_{\eta} \eta^T z - (H(\beta^* + \eta) - H(\beta^*)) \right\} = - \inf_{z \in \mathcal{K}} \left\{ \sup_{\eta} \eta^T z - \Lambda(\eta; \beta^*) \right\}. \quad (39)$$

To prove Theorem 4.1, apply (39) to a Gaussian random variable with mean $X^*\beta^*$ and covariance matrix $\sigma^2 I_n$ to get the rate function as:

$$\rho(z) = \inf_{\eta} -\eta^T(z - X^*\beta^*) + \frac{1}{2}\sigma^2\eta^T\eta,$$

by plugging log-MGF of Y evaluated at η as $\Lambda(\eta; \beta^*) = \eta^T X^*\beta^* + \frac{1}{2}\sigma^2\eta^T\eta$. The above infimum occurs at $\eta = \frac{(z - X^*\beta^*)}{\sigma^2}$, yielding the bound

$$\sup_{z \in \mathcal{K}} -\frac{1}{2\sigma^2} \|(z - X^*\beta^*)\|^2 = - \inf_{z \in \mathcal{K}} \frac{1}{2\sigma^2} \|(z - X^*\beta^*)\|^2,$$

which proves

$$\log \mathbb{P}(Y \in \mathcal{K} | \beta^*) \leq - \frac{\|X^*\beta^* - \mathcal{P}_{\mathcal{K}}(X^*\beta^*)\|^2}{2\sigma^2}.$$

□

Proof of Theorem A.1. To prove (39),

$$\log \mathbb{P}(Z \in \mathcal{K}) = \log \mathbb{E}(1_{\{Z \in \mathcal{K}\}}) \leq \log \mathbb{E}(\exp(\eta^T Z + u))$$

for any $\eta \in \mathbb{R}^d$ and for any u satisfying $\eta^T z + u \geq 0 \forall z \in \mathcal{K}$.

Upon fixing a choice of u as $u_{\text{opt}}(\eta) = \sup_{z \in \mathcal{K}} -\eta^T z$, the upper bound can be re-written as:

$$\begin{aligned} \log \mathbb{P}(Z \in \mathcal{K}) &\leq u_{\text{opt}} + \log \mathbb{E}(\exp(\eta^T Z) | \beta^*) \text{ for all } \eta \in \mathbb{R}^d \\ &= \sup_{z \in \mathcal{K}} -\eta^T z + (H(\beta^* + \eta) - H(\beta^*)) \text{ for all } \eta \in \mathbb{R}^d \\ &\leq \inf_{\eta} \left\{ \sup_{z \in \mathcal{K}} -\eta^T z + (H(\beta^* + \eta) - H(\beta^*)) \right\} \\ &= - \sup_{\eta} \left\{ \inf_{z \in \mathcal{K}} \eta^T z - (H(\beta^* + \eta) - H(\beta^*)) \right\} \\ &= - \inf_{z \in \mathcal{K}} \left\{ \sup_{\eta} \eta^T z - (H(\beta^* + \eta) - H(\beta^*)) \right\} \end{aligned}$$

$$= - \inf_{z \in K} \left\{ \sup_{\eta} \eta^T z - \Lambda(\eta; \beta^*) \right\}.$$

The third inequality follows by an optimization over $\eta \in \mathbb{R}^d$ and the last one follows from a minimax theorem argument by noting that the function over which the optimization takes place, given by

$$f(z, \eta) = \eta^T z - (H(\beta^* + \eta) - H(\beta^*)).$$

is convex-concave and continuous as a map of z for every $\eta \in \mathbb{R}^d$. \square

Proof of Lemma 6.3. Differentiating (23) for $w(z) = w^{\text{bar}}(z)$ with respect to β^* gives the estimating equation

$$X^{*T} \nabla \bar{w}^*(X^* \hat{\beta}^*) = X^{*T} y + \nabla \log \pi(\hat{\beta}^*).$$

Now, observing that the maximizer of objective

$$z^T \mu - \frac{1}{2} \|z\|^2 - \sum_{j=1}^m \log(1 + 1/(b_j - a_j^T z)),$$

satisfies $\nabla \bar{w}(z^*(\mu)) = \mu$, the statement follows. \square

Proof of Lemma 6.4. The randomized selective MAP satisfies

$$-y + \frac{\gamma^2}{1 + \gamma^2} \hat{\beta}^* + \frac{1}{(1 + \gamma^2)} \nabla \bar{w}^* \left(\frac{\hat{\beta}^*}{1 + \gamma^2} \right) - \nabla \log \pi(\hat{\beta}^*) = 0,$$

where we note that the selection probability is approximated by the optimization problem

$$\bar{w}^*(\mu) = \sup_{z \in \mathbb{R}^n} z^T \mu - \bar{w}(z),$$

with $w(z)$ defined for the polyhedral region $Ay \leq b$. This implies that

$$\nabla \bar{w}^* \left(\frac{\hat{\beta}^*}{1 + \gamma^2} \right) = (1 + \gamma^2)(y + \nabla \log \pi(\hat{\beta}^*)) - \gamma^2 \hat{\beta}^*,$$

which results in estimating equation

$$\hat{\beta}^* = (1 + \gamma^2) \nabla \bar{w} \left((1 + \gamma^2)(y + \nabla \log \pi(\hat{\beta}^*)) - \gamma^2 \hat{\beta}^* \right).$$

\square

Proof of Theorem 7.1. It suffices to show

$$-\inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*) \leq \liminf_n \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) \leq \limsup_n \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) \leq -\inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*)$$

in order to prove 29. To see the upper bound, we have along similar lines of theorem A.1 in the paper that

$$\begin{aligned} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) &\leq \log \mathbb{E}(\exp(n\eta^T \bar{Y}_n + u) | \beta^*) \text{ for every } \eta \text{ and } u : n\eta^T x + u \geq 0 \text{ for } x \in \mathcal{K} \\ &\leq n \cdot \sup_{x \in \mathcal{K}} -\eta^T x + \Lambda(\eta; \beta^*) \\ &\leq -n \cdot \sup_{\eta} \{ \inf_{x \in \mathcal{K}} \eta^T x - \Lambda(\eta; \beta^*) \} \\ &= -n \cdot \inf_{x \in \mathcal{K}} \{ \sup_{\eta} \eta^T x - \Lambda(\eta; \beta^*) \} \text{ by a minimax argument} \\ &= -n \inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*). \end{aligned}$$

This completes the proof that

$$\limsup_n \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) \leq -\inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*).$$

To prove the lower bound, it suffices to show for every $x \in \mathcal{K}^o$ and any $\delta > 0$ that

$$\liminf_n \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{B}(x, \delta) | \beta^*) \geq -\Lambda^*(x; \beta^*) \quad (40)$$

where $\mathcal{B}(\bar{x}, \delta) = \{y : \|y - x\|_\infty \leq \delta\}$. This shall prove

$$\liminf_n \frac{1}{n} \log \mathbb{P}(\bar{Y}_n \in \mathcal{K} | \beta^*) \geq -\inf_{x \in \mathcal{K}^o} \Lambda^*(x; \beta^*) = -\inf_{x \in \mathcal{K}} \Lambda^*(x; \beta^*). \quad (41)$$

For this $x \in \mathcal{K}^o$, denote

$$\eta^* = \operatorname{argsup}_{\eta} \eta^T x - \Lambda(\eta; \beta^*)$$

which satisfies $\nabla \Lambda(\eta^*; \beta^*) = x$. To prove the above, define an exponentially tilted measure $\tilde{\mathbb{P}}$ as

$$d\tilde{\mathbb{P}}(z) = \exp(\eta^{*T} z - \Lambda(\eta^*)) d\mathbb{P}(z).$$

For $0 < \delta_1 < \delta$

$$\begin{aligned} \mathbb{P}(\bar{Y}_n \in \mathcal{B}(x, \delta)) &\geq \mathbb{P}(\bar{Y}_n \in \mathcal{B}(x, \delta_1)) \\ &= \exp(n\Lambda(\eta^*; \beta^*)) \int_{\mathcal{B}(\bar{z}, \delta_1)} \exp(-n\eta^{*T} x) \exp(-\eta^{*T} \{(z_1 + z_2 + \dots + z_n) - nx\}) d\tilde{\mathbb{P}}(z_1) \dots d\tilde{\mathbb{P}}(z_n) \\ &= \exp(-n\Lambda^*(x; \beta^*) - n|\eta^*|^T \delta_1) \end{aligned}$$

Finally, we have (40) by letting $\delta_1 \rightarrow 0$ in the above lower bound. \square

Proof of Corollary 7.2. It follows from 7.1 that the sequence of true selective posteriors, written as

$$\log \pi_S(\beta^*|\bar{y}_n) = \log \pi(\beta^*) - n\|\bar{y}_n - \beta^*\|^2/2 - \log \mathbb{P}(\bar{Y}_n \in \mathcal{K}|\beta^*)$$

can be approximated by

$$\log \pi(\beta^*) - n\|\bar{y}_n - \beta^*\|^2/2 + n \inf_{x \in \mathbb{R}^p} \{\|x - \beta^*\|^2/2\}.$$

Now, it is easy to see that as long as the sequence of convex barrier functions $\frac{1}{n}w_{1/\sqrt{n}}(x) \rightarrow I_{\mathcal{K}}(x)$ for all $x \in \mathbb{R}^d$ as $n \rightarrow \infty$, the limiting sequence of objectives

$$\|x - \beta^*\|^2/2 + \frac{1}{n}w_{1/\sqrt{n}}(x)$$

are convex in x and converge to the continuous, convex objective $\|x - \beta^*\|^2/2 + I_{\mathcal{K}}(x)$, which has a unique minimum. Hence, we also have

$$\inf_{x \in \mathbb{R}^d} \left\{ \|x - \beta^*\|^2/2 + \frac{1}{n}w_{1/\sqrt{n}}(x) \right\} \rightarrow \inf_{x \in \mathbb{R}^d} \left\{ \|x - \beta^*\|^2/2 + I_{\mathcal{K}}(x) \right\} \text{ as } n \rightarrow \infty.$$

□

Proof of Theorem 7.3. By Theorem 6.2, the maximizer of (30) is

$$\hat{\beta}_n^* = \bar{Y}_n - \frac{1}{\sqrt{n}(\sqrt{n}\bar{Y}_n)\sqrt{n}\bar{Y}_n + 1}.$$

Denoting $Z_n = (\sqrt{n}|\beta^*|)\sqrt{n}\bar{Y}_n$ and $b(z) = \log\left(1 + \frac{1}{z}\right)$ we have

$$\begin{aligned} \hat{\beta}_n^* - \beta^* &= \bar{Y}_n - \beta^* - n^{-1/2} \frac{1}{\sqrt{n}\bar{Y}_n\sqrt{n}\bar{Y}_n + 1} \\ &= \bar{Y}_n - \beta^* + n^{-1/2} \nabla b(n^{1/2}\bar{Y}_n) \\ &= \frac{Z_n}{n|\beta^*|} - \beta^* + n^{-1/2} \nabla b\left(\frac{Z_n}{n^{1/2}|\beta^*|}\right) \end{aligned}$$

For $\beta^* = -\delta < 0$, $Z_n = (n^{1/2}|\beta^*|)n^{1/2}\bar{Y}$ converges in law to an exponential random variable with mean 1. Further note that,

$$z\nabla b(z) \rightarrow -K \text{ as } z \downarrow 0$$

for $K = 1$. Using these two facts, we have $\frac{Z_n}{n|\beta^*|} = o_p(1)$ and

$$n^{-1/2} \nabla b\left(\frac{Z_n}{n^{1/2}|\beta^*|}\right) = \left(\frac{|\beta^*|}{Z_n}\right) \frac{Z_n}{n^{1/2}|\beta^*|} \nabla b\left(\frac{Z_n}{n^{1/2}|\beta^*|}\right) = \left(\frac{|\beta^*|}{Z_n}\right) W_n$$

where $W_n \rightarrow -K$ as $n \rightarrow \infty$. Hence, we can approximate the sequence of random variables $\hat{\beta}_n^* - \beta^*$ in distribution by the random variable

$$-\frac{K|\beta^*|}{Z} - \beta^* \text{ where } Z \sim \text{Exp}(1).$$

Set $\beta^* = -\delta$ for $\delta > 0$

$$\begin{aligned} \mathbb{P}\left(|\hat{\beta}_n^* - \beta^*| > \frac{\delta}{2}|\beta^*|\right) &\approx \mathbb{P}\left(\left|\frac{K|\beta^*|}{Z} + \beta^*\right| > \frac{\delta}{2}|\beta^*|\right) \\ &\geq \mathbb{P}\left(-\frac{K|\beta^*|}{Z} - \beta^* > \frac{\delta}{2}|\beta^*|\right) \\ &= \mathbb{P}\left(Z > -\frac{K|\beta^*|}{\beta^* + \delta/2} \mid \beta^*\right) \\ &= \exp\left(-\frac{K|\beta^*|}{\beta^* + \delta/2}\right) = \exp(2K) > 0. \end{aligned}$$

□

Proof of Lemma 7.4. $\hat{C}_n(\cdot)$ is a strongly convex as $\frac{\gamma^2}{2(1+\gamma^2)}\|\beta^*\|^2$ is strongly convex with index $\gamma^2/(1+\gamma^2)$ and $\bar{w}_n^*(\beta^*)$ is a convex function in β^* . It is straight forward from here to see that the indices of convexity of the sequence of approximate log partition functions are bounded below by $M = \gamma^2/(1+\gamma^2) > 0$. The same goes for the sequence $-\ell_S^n(\beta^*)$ which has a linear term $-\beta^* \bar{y}_n$ added to $\hat{C}_n(\beta^*)$. □

Proof of Lemma 7.5. The randomized selective MLE $\hat{\beta}_n^*$ satisfies

$$\nabla \hat{C}_n(\hat{\beta}_n^*) = \bar{y}_n, \text{ that is } \hat{\beta}_n^* = \nabla \hat{C}_n^{-1}(\bar{y}_n).$$

Thus, we have

$$\|\hat{\beta}_n^* - \beta^*\|^2 = \|\nabla \hat{C}_n^{-1}(\bar{y}_n) - \beta^*\|^2 = \|\nabla \hat{C}_n^*(\bar{y}_n) - \nabla \hat{C}_n^*(\nabla \hat{C}_n(\beta^*))\|^2.$$

Lemma 7.4 shows that $\hat{C}_n(\beta^*)$ is strongly convex with indices of convexity $m_n \geq M = \gamma^2/(1+\gamma^2)$ and the proof is complete by using the fact that the convex conjugate of a strongly convex function with index m_n is Lipschitz smooth with Lipschitz index $1/m_n$. Hence, we have

$$\|\hat{\beta}_n^* - \beta^*\|^2 \leq \frac{1}{m_n^2} \|\bar{y}_n - \nabla \hat{C}_n(\beta^*)\|^2 \leq \frac{1}{M^2} \|\bar{y}_n - \nabla \hat{C}_n(\beta^*)\|^2.$$

□

Proof of Theorem 7.6. For a fixed $\epsilon > 0$, Markov's inequality and 7.5 yields

$$\mathbb{P}(\|\hat{\beta}_n^* - \beta^*\| > \epsilon \mid \bar{Y}_n^* \in \mathcal{K}) \leq \frac{\mathbb{E}[n\|\hat{\beta}_n^* - \beta^*\|^2 \mid \bar{Y}_n^* \in \mathcal{K}]}{n\epsilon^2} \leq \frac{\mathbb{E}[n\|\bar{Y}_n - \nabla \hat{C}_n(\beta^*)\|^2 \mid \bar{Y}_n^* \in \mathcal{K}]}{n\epsilon^2}.$$

Denoting $C_n(\beta^*)$ as the log partition function of the true randomized selective likelihood $f_S^n(\beta^*)$, we have

$$\begin{aligned} \frac{\mathbb{E}[n\|\bar{Y}_n - \nabla \hat{C}_n(\beta^*)\|^2 \mid \bar{Y}_n^* \in \mathcal{K}]}{n\epsilon^2} &= \frac{\mathbb{E}[n\|\bar{Y}_n - \nabla C_n(\beta^*)\|^2 \mid \bar{Y}_n^* \in \mathcal{K}]}{n\epsilon^2} + \frac{\|\nabla C_n(\beta^*) - \nabla \hat{C}_n(\beta^*)\|^2}{\epsilon^2} \\ &= \frac{\text{Var}[\sqrt{n}\bar{Y}_n \mid \bar{Y}_n^* \in \mathcal{K}]}{n\epsilon^2} + \frac{\|\nabla C_n(\beta^*) - \nabla \hat{C}_n(\beta^*)\|^2}{\epsilon^2} \\ &\leq \frac{\text{Var}[\sqrt{n}\bar{Y}_n]}{n\epsilon^2} + \frac{\|\nabla C_n(\beta^*) - \nabla \hat{C}_n(\beta^*)\|^2}{\epsilon^2} \end{aligned}$$

The last step uses the fact that variance of a Gaussian random variable reduces when restricted to a convex set (see Kanter and Proppe (1977) for a proof) and thus, the first term converges to 0 as $n \rightarrow \infty$. Convergence of the second term to 0 follows from a combination of corollary 7.2 and the properties of convexity and differentiability $\hat{C}_n(\beta^*)$. \square

Proof of Theorem 7.7. We shall show that

$$\mathbb{P}\left(R_n(\delta) \leq \frac{1}{2} \inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\}\right) \leq \mathbb{P}(\|\hat{\beta}_n^* - \tilde{\beta}_n^*\| \leq \delta)$$

where $R_n(\delta) = \sup_{\|s - \hat{\beta}_n^*\| \leq \delta} |\ell_S^n(s) - \hat{\ell}_S^n(s)|$. Let $p_n = \hat{\beta}_n^* + \alpha u$ for a unit vector u and $\alpha > \delta$. Then

$$\begin{aligned} \ell_S^n(\hat{\beta}_n^* + \delta u) &= \ell_S^n\left(\left(1 - \frac{\delta}{\alpha}\right)\hat{\beta}_n^* + \frac{\delta}{\alpha}p_n\right) \geq \left(1 - \frac{\delta}{\alpha}\right)\ell_S^n(\hat{\beta}_n^*) + \frac{\delta}{\alpha}\ell_S^n(p_n). \\ \frac{\delta}{\alpha}\{\ell_S^n(\hat{\beta}_n^*) - \ell_S^n(p_n)\} &\geq \ell_S^n(\hat{\beta}_n^*) - \ell_S^n(\hat{\beta}_n^* + \delta u) \\ &= (\ell_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(\hat{\beta}_n^*)) - \{\ell_S^n(\hat{\beta}_n^* + \delta u) - \hat{\ell}_S^n(\hat{\beta}_n^* + \delta u)\} + (\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(\hat{\beta}_n^* + \delta u)) \\ &\geq \inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\} - 2 \cdot \sup_{\|s - \hat{\beta}_n^*\| \leq \delta} |\ell_S^n(s) - \hat{\ell}_S^n(s)| \end{aligned}$$

We note that the event

$$R_n(\delta) \leq \frac{1}{2} \inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\},$$

implies $\ell_S^n(\hat{\beta}_n^*) - \ell_S^n(p_n) > 0$ for all $p_n = \hat{\beta}_n^* + \alpha u$ and for any $\alpha > \delta$, which means that the maximizer of $\ell_S^n(\cdot)$ given by $\tilde{\beta}_n^*$ lies inside a δ ball around $\hat{\beta}_n^*$, the maximizer of the pseudo selective posterior sequence. Thus, we have

$$\mathbb{P}(\|\hat{\beta}_n^*(\bar{y}_n) - \tilde{\beta}_n^*(\bar{y}_n)\| \geq \delta) \leq \mathbb{P}\left(R_n(\delta) \geq \frac{1}{2} \inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\}\right).$$

The above equation holds for any choice of norm $\|\cdot\|$. To complete the proof, we note from 7.4 that negative of $\hat{\ell}_S^n$ is strongly convex with index m_n and hence, for any s

$$\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s) \geq \frac{1}{2}m_n\|s - \hat{\beta}_n^*\|^2,$$

which implies

$$\inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\} \geq \frac{1}{2}m_n\delta^2 \geq \frac{1}{2}M\delta^2.$$

Using the fact that the randomized selective MLE $\hat{\beta}_n^*$ is stochastically bounded under the selective law (follows from Theorem 7.6) and uniform convergence of $R_n(\delta)$ on compact sets (from Remark 6)

$$\mathbb{P}(\|\hat{\beta}_n^*(\bar{y}_n) - \tilde{\beta}_n^*(\bar{y}_n)\| \geq \delta) \leq \mathbb{P}\left(R_n(\delta) \geq \frac{1}{2} \inf_{s: \|s - \hat{\beta}_n^*\| = \delta} \{\hat{\ell}_S^n(\hat{\beta}_n^*) - \hat{\ell}_S^n(s)\}\right) \geq \mathbb{P}\left(R_n(\delta) \geq \frac{1}{2}M\delta^2\right) \rightarrow 0.$$

□

B An Algorithm for Computing $\mathcal{P}_K(\mu)$

We present a numerical optimization algorithm that can be implemented to find the projection of mean vector μ onto the convex selection region $\mathcal{K} = \{y : A_E y \leq b_E\}$ for approximating the Gaussian volume of selection region \mathcal{K} . The convex problem of interest can be framed as below

$$\min_{\mu: A_E \mu \leq b_E} \frac{1}{2} \|(\mu - X_E \beta_E)\|^2.$$

We employ the ADMM algorithm (Boyd et al., 2011), which augments the term $\frac{\rho}{2} \|(A_E \mu - b_E - r)\|^2$ to the Lagrangian to solve the optimization problem:

$$\min_{\mu, r, y: r = A_E \mu - b_E} \frac{1}{2} \|(\mu - X_E \beta_E)\|^2 + y^T (A_E \mu - b_E - r) + \frac{\rho}{2} \|(A_E \mu - b_E - r)\|^2 + \mathbb{1}_{r \leq 0}$$

The updates for the ADMM algorithm are given by:

$$\mu^{(k+1)} = (I + \rho A_E^T A_E)^{-1} (X_E \beta_E + \rho A_E^T (b_E + r^{(k)} - u^{(k)})) \quad (42)$$

$$r^{(k+1)} = \min (A_E \mu^{(k+1)} - b_E + u^{(k)}, 0) \quad (43)$$

$$u^{(k+1)} = u^{(k)} + (A_E \mu^{(k+1)} - b_E - r^{(k+1)}) \quad (44)$$

We can adopt the adaptive Metropolis-Hastings with the selection probability for each fresh draw computed via the above the ADMM updates.